

DB2 Planning

Presented by DB2 Developer Domain

<http://www7b.software.ibm.com/dmdd/>

Table of Contents

If you're viewing this document online, you can click any of the topics below to link directly to that section.

1. Introduction	2
2. DB2 products	4
3. DB2 tools	12
4. The Control Center	23
5. The Configuration Assistant	35
6. Other DB2 tools	37
7. Data warehousing	44
8. OLAP	51
9. Resources and feedback	53

Section 1. Introduction

What this tutorial is about

This tutorial introduces the basics of the DB2 products and tools, along with the concepts of data warehousing and OLAP. It is the first in a series of six tutorials designed to help you prepare for the DB2 UDB V8.1 Family Fundamentals Certification (Exam 700). The material in this tutorial primarily covers the objectives in Section 1 of the test, which is entitled "Planning." You can view these objectives at:

<http://www.ibm.com/certify/tests/obj700.shtml>

DB2 installation is not covered in this tutorial. If you haven't already done so, we strongly recommend that you download and install a copy of *IBM DB2 Universal Database*, Enterprise Server Edition. Installing DB2 will help you understand many of the concepts that are tested on the DB2 UDB V8.1 Family Fundamentals Certification exam. The installation process is documented in the Quick Beginnings books, which can be found at the *DB2 Technical Support* Web site under the Technical Information heading.

About the author

Paul C. Zikopoulos, BA, MBA, is an IDUG keynote and award-winning speaker with the IBM Global Sales Support team. He has more than seven years of experience with DB2 and has written numerous magazine articles and books about it. Paul has co-authored several books: *DB2: The Complete Reference*, *DB2 Fundamentals Certification for Dummies*, *DB2 For Dummies*, and *A DBA's Guide to Databases on Linux*. Paul is a DB2 Certified Advanced Technical Expert (DRDA and Cluster/EEE) and a DB2 Certified Solutions Expert (Business Intelligence and Database Administration). You can reach him at paulz_ibm@msn.com.

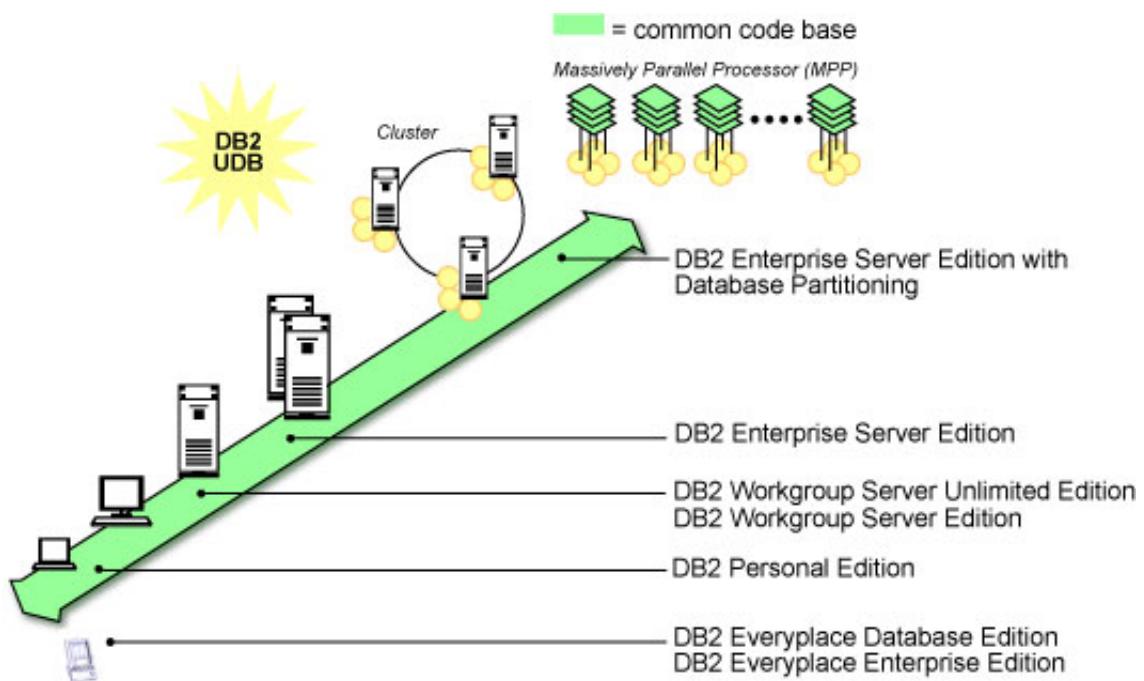


Section 2. DB2 products

The different editions of DB2

DB2 Version 8 delivers data management solutions for any business. No other database management system can match the advanced performance, availability, scalability, and manageability features that DB2 Version 8 provides. It is easy to arrive at the decision to use DB2. However, there are different editions of DB2 available, each suited to a different part of the marketplace. On the Fundamentals exam you are expected to understand the different DB2 products that are available, so I will cover them in this section.

All the available distributed editions of DB2 are shown in the figure below. The figure represents a progression: each edition displayed includes all the functions, features, and benefits of the editions to its left, along with new features and functionality. (DB2 Everyplace is an exception to this rule, as we'll see.)



DB2 Everyplace

The true power of mobile computing lies not in the mobile device itself, but in its ability to tap into data from other sources. DB2 Everyplace brings the power of DB2 to mobile devices, leveraging their ability to synchronize data with other systems -- literally putting your enterprise data in the pockets of your mobile workforce and giving them the ability to update your enterprise data from the field.

DB2 Everyplace is more than just a mobile computing infrastructure. It's a complete development environment that includes the tools you need to build, deploy, and support powerful e-business applications. DB2 Everyplace features a tiny "fingerprint" engine (about 180 KB) packed full of security features, like table encryption, and advanced indexing techniques that lead to high performance. It can comfortably run (with multithreaded support) on a wide variety of handheld devices, with support for Palm OS, Microsoft Windows CE/Pocket PC, any Microsoft Windows 32-bit operating system, Symbian, QNX Neutrino, Java 2 Platform Micro Edition (J2ME) devices, and embedded Linux distributions (such as BlueCat Linux).

If you need a relational engine (and synchronization services) on a constrained device, you must use DB2 Everyplace. You should also consider this product for occasionally connected mobile users on laptops if their applications don't need features (like triggers) that are not part of the DB2 Everyplace engine.

DB2 Personal Edition

DB2 Personal Edition (DB2 PE) is a single-user RDBMS that runs on low-cost commodity hardware desktops. DB2 PE is available for Windows 98, Windows ME, Windows NT (SP6 or later), Windows 2000 (SP2 recommended), Windows XP, and Linux. Windows 2003 servers will be supported when Microsoft releases them. DB2 PE has all of the features of DB2 Workgroup Server Edition, with one exception: remote clients cannot connect to databases that are running this edition of DB2. (However, workstations with the Control Center can connect to these databases to perform remote administration.) Because "DB2 is DB2 is DB2," applications that are developed for DB2 PE will run on any other edition of DB2. You can use DB2 PE to develop DB2 applications before rolling them out into a production environment.

DB2 PE is useful both for PCs that are not connected to a network and for those that are; in either case, it is useful for users who need a powerful data store, or who need to provide database storage facilities and be able to connect to remote DB2 servers. DB2 PE also comes with the DB2 Warehouse Center, so users can experiment with business intelligence modeling functions like extract, transform, and load (see [What is data warehousing? on page 44](#)). DB2 PE is extensible with support for all the DB2 Extenders: XML (now bundled into a DB2 installation), Spatial, and Net Search. (See [DB2 Extenders](#) on page 8 for more information.)

Occasionally connected users may want to take advantage of DB2's built-in replication feature and the DB2 Control Server to set up a synchronized environment where mobile workers can keep in touch with their enterprise. Of course, this would only be suitable for users of laptops and certain workstations, such as those running point-of-sale (POS) applications.

DB2 Workgroup Server Edition

DB2 Workgroup Server Edition (DB2 WSE) is a full-function, Web-enabled client/server

RDBMS. It is available on all supported flavors of UNIX (AIX, Solaris, and HP-UX), Linux, Windows NT (SP6 or later), Windows 2000 (SP2 recommended), and Windows XP. DB2 WSE will support the Windows .NET environment when Microsoft releases it.

DB2 WSE provides a low-cost, entry-level server that is intended primarily for small business and departmental computing. For the most part, it is functionally equivalent to DB2 Enterprise Server Edition, though there is no integrated mainframe connectivity via the DB2 Connect component, and there are some limits to its extensibility and features (for example, DB2 WSE does not support 64-bit computing or DB2 Data Links File Manager). As previously noted, an application written for any edition of DB2 is transparently portable to another edition on any distributed platform.

In UNIX, Linux, and Windows environments, you can install DB2 WSE on symmetric multiprocessing (SMP) machines with up to four processors. DB2 WSE is licensed by the client-to-server model using a *concurrent* (sometimes referred to as *capacity pricing*) or registered user option.

DB2 WSE can play many roles in a business. It is a good fit for small or medium-sized businesses (SMBs) that need a full-fledged relational database store that is scalable and available over a wide area network (WAN) or local area network (LAN). It is also useful for enterprise environments that need silo servers for lines of business, or for departments that need the ability to scale in the future. DB2 WSE's licensing options do not allow for external Web-based usage; if you want an SMB-targeted edition of DB2 for use over the Web, you should consider DB2 Workgroup Unlimited Server Edition (DB2 WSUE). DB@ WSE can be used internally for Web-access behind a company's firewall. In this case, you must be able to count each user that is accessing the DB2 server over the company's intranet.

Note: In DB2 Version 7, this product was called *DB2 Workgroup Edition*.

DB2 Workgroup Server Unlimited Edition

DB2 Workgroup Server Unlimited Edition (DB2 WSUE) is essentially DB2 WSE with different licensing terms. For a complete description of the functionality available with this product, see the description of DB2 WSE on the previous panel.

DB2 WSUE can only be licensed through the processor option and is the right choice for SMBs that want to use the Web to transact their business or to streamline their processes.

Because this DB2 WSUE is accompanied by a processor license, you can use it to serve up Internet *and* Intranet requests that come to the database manager via an application server. This licensing model eliminates the need to count users accessing the database. DB2 WSUE can support up to four-way SMP servers on UNIX- or Intel-based equipment.

Users will typically look at this edition of DB2 if they need to serve up DB2 data over both the Internet and an Intranet, but do not need the added functions, features, and benefits of DB2 Enterprise Server Edition (DB2 ESE). DB2 WSUE cannot create 64-bit

instances, does not come with the connection concentrator, and does not include the integrated DB2 Connect component found in DB2 ESE.

Note: In DB2 Version 7, this product was called *DB2 Workgroup Unlimited Edition*.

DB2 Enterprise Server Edition

DB2 Enterprise Server Edition (DB2 ESE) is a full-function, Web-enabled client/server RDBMS. It is available on all supported flavors of UNIX (AIX, Solaris, and HP-UX), Linux, Windows NT (SP6 or later), and Windows 2000 (SP2 recommended). Version 8 of DB2 ESE does not run on Windows XP for production purposes, but can be licensed for user acceptance testing, test, and application development on that OS. This is not a DB2 limitation: Microsoft limits the number of concurrent connections that can be handled by the single-user version of their Windows server products. The concurrent connection limits can be handled easily by a four-way SMP box and therefore DB2 WSE or DB2 WSUE would be a better choice (unless you need some of the features that are only included in DB2 ESE). DB2 ESE will support the Windows .NET environment when Microsoft releases it.

DB2 ESE is meant for large and mid-sized departmental servers. DB2 ESE includes all the functionality provided by DB2 WSE; in addition, it includes the DB2 Connect component, which enables you to connect to iSeries- and zSeries-based DB2 databases, as well as non-database host resources like CICS, VSAM, and IMS. In DB2 Version 8, each server running DB2 ESE is licensed for five registered users when connecting to host-based data sources. If you need additional connectivity, you should purchase a separate DB2 Connect server license and user entitlements.

DB2 ESE is licensed per processor and can be installed on any number of processors on a single SMP box, or on multiple boxes. You can use this edition of DB2 for any applications.

DB2 ESE has the ability to partition data within a single server, across multiple database servers (all of which have to be running on the same operating system), or within a large SMP machine out of the box, thanks to its database partitioning feature (DPF). You can purchase DPF as part of a DB2 ESE processor license. With DPF, the size of your database is only limited by the number of computers you have. DB2 ESE with DPF is meant for larger data warehouses, or for high-performance online transaction processing (OLTP) requirements. DB2 ESE with the DPF also allows multiple SMP machines to be clustered together under a single database image for very large-scale transaction volumes.

Note: In DB2 Version 7, this product was called *DB2 Enterprise Edition*. The database partitioning feature that allows for the creation of a partitioned database was sold as *DB2 Enterprise Extended Edition*. Because these two DB2 products always had the same codebase, in DB2 Version 8 they have merged to become DB2 ESE with the DPF.

DB2 Universal Developer's and DB2 Personal Developer's Edition

A special offering called DB2 Universal Developer's Edition (DB2 UDE) is available for application developers. This is a reduced-price offering that gives application developers access to the DB2 Extenders and all distributed flavors of DB2. This package ships copies of DB2 ESE that are restricted to the development, evaluation, demonstration, and testing of application programs.

DB2 Personal Developer's Edition ships with all the personal DB2 products and is available free of charge (see [What this tutorial is about](#) on page 2).

DB2 clients

DB2 ships clients that are used to communicate with DB2 servers. There are three kinds of DB2 clients:

- **DB2 Runtime Client:** This client provides the ability for workstations running a variety of platforms to access DB2 databases. It provides just basic connectivity -- nothing more and nothing less. If you need to establish connectivity to a remote DB2 server or DB2 Connect Gateway (which helps you access DB2 on a mainframe or host system like DB2 for z/OS), you have to start here at minimum. Of course, you could use any client for connectivity.
 - **DB2 Administration Client:** This client provides the ability for workstations from a variety of platforms to access and administer DB2 databases through the Control Center or the Configuration Assistant. A DB2 Administration client has all the features of a DB2 Runtime client, and also includes all of the DB2 Administration tools, documentation, and support for thin clients. It also provides the ability (if you have DB2 Connect) to manage z/OS and IMS subsystems.
 - **DB2 Application Development Client:** This client provides the tools and environment you need to develop applications that access DB2 servers. You can build and run DB2 applications with a DB2 Application Development client. Of course, since this is a DB2 client, it gives users the power of connectivity as well.
-

DB2 Extenders

The DB2 Extenders can take your database applications beyond traditional numeric and character data to images, XML, videos, voice, spatial objects, complex documents, and more. Using extenders, you can bring all these types of data into a database and

work with them using SQL -- the language used to talk to relational databases. (For more detailed information on SQL, see the [fourth tutorial in this series](#).) Imagine, selecting a set of wallpaper styles based on the color or patterns you are looking to find: that's extender technology!

On this panel, we'll outline most of the DB2 Extenders. You can find out more about all of the DB2 Extenders at the DB2 Extenders Web site (see [Resources](#) on page 53).

- **XML Extender:** DB2's XML Extender provides new data types that let you store XML documents in DB2 databases, and adds functions that help you work with these XML documents while in a database.

You can store entire XML documents in DB2, or store them as external files managed by the database. This method is known as *XML Columns*. You can also decompose an XML document into relational tables and then recompose that information to XML on the way out of the database. Basically, this means that your DB2 database can strip the XML out of a document and just take the data, or take data and create an XML document from it. This method is known as *XML Collections*.

In DB2 Version 7, the XML Extender was a free, separately installable product. In DB2 Version 8, it is built into the DB2 installation as a component.

- **DB2 Net Search Extender:** This extender helps businesses that need fast performance when searching for information in a database. You are likely to see this used in Internet applications, where excellent search performance on large indexes and scalability of concurrent queries are needed. So, if you need a high-speed in-memory search, this is the extender for you.

In DB2 Version 8, the Text Information Extender has merged with the Net Search Extender. This merge adds smart, usable text search extensions to your database and provides you with more intelligent information related to your business needs.

High-performance in-memory searches are indispensable for e-commerce applications or any other application with high performance and scalability text-search demands. For example, you can now use in-memory search with your XML documents. A license for up to five concurrent users comes with DB2 WSE with this product. If you want to use this product and support more than five users, you must purchase it and run it on DB2 WSUE or DB2 ESE.

- **DB2 Spatial Extender:** This extender allows you to store, manage, and analyze *spatial data* -- that is, information about the location of geographic features -- in DB2, along with traditional data for text and numbers. With this capability, you can generate, analyze, and exploit spatial information about geographic features, such as the locations of office buildings or the size of a flood zone. The DB2 Spatial Extender extends the function of DB2 with a set of advanced spatial data types that represent geometries such as points, lines, and polygons; it also includes many functions and features that interoperate with those new data types. These capabilities allow you to integrate spatial information with your business data, adding another element of intelligence to your database. A license for up to five concurrent users comes with DB2 WSE with this product. If you want to use this product and support more than five users, you must purchase it and run it on DB2 WSUE or DB2

ESE.

- **Text, Audio, Image, and Video (TAIV) Extenders:** These extenders allow you to extend the relational database to use nontraditional forms of data like text, songs, pictures, and movies. With the TAIV Extenders, you can work with this data via SQL.
-

DB2 Data Links File Manager

Data Links is cool IBM technology that represents the next generation of enterprise content management. This is an add-on product to DB2. Data Links allows you to manage files that reside outside the database as though they were logically within the database. Data Links guarantees referential integrity to these external files, provides enhanced access control to them, and supports the automatic and coordinated backup and restore capability within transactional environments that are crucial for data management. It simplifies and reduces system administration costs and complexities by providing a single coordinated administration point for file and database data. See [Resources](#) on page 53 for more information on DB2 Data Links Manager.

DB2 Connect

A great deal of the data in many large organizations is managed by DB2 for AS/400, DB2 for MVS/ESA, DB2 for z/OS, or DB2 for VSE and VM. Applications that run on any of the supported DB2 distributed platforms can work with this data transparently, as if a local database server managed it. You can also use a wide range of off-the-shelf or custom-developed database applications with DB2 Connect and its associated tools.

DB2 Connect provides connectivity to mainframe and midrange databases from Windows, Linux, and UNIX platforms.

There are a number of DB2 Connect editions available: Personal Edition, Enterprise Edition, Application Server Edition, and Unlimited Edition. This product is an add-on product to DB2 that can be purchased separately, although some complimentary user licenses are provided in DB2 ESE. See [Resources](#) on page 53 for more information on DB2 Connect.

DB2 add-on tools

There are two kinds of tools for DB2: those that are free and those that are add-ons that can be purchased separately. The free tools come as part of a DB2 installation and can be launched from the Control Center, the Configuration Assistant, or on their own (you will learn about them later in this tutorial). A separate set of tools is available to

help ease the DBA's task of managing and recovering data and making it accessible.

There are five tools that can be purchased for distributed versions of DB2:

- **DB2 Recovery Expert:** Provides simplified, comprehensive, and automated recovery with extensive diagnostics and SMART (self-managing and resource tuning) techniques to minimize outage duration.
- **DB2 Performance Expert:** Integrates performance monitoring, reporting, buffer pool analysis, and a performance warehouse function into one tool. It provides a *single* system overview that monitors all subsystems and instances across many different platforms in a consistent way.
- **DB2 High Performance Unload:** Quickly and efficiently unloads and extracts data from DB2 for movement across enterprise systems.
- **DB2 Web Query Tool:** Connects all your users directly to multiple enterprise databases, securely and simultaneously, regardless of database size, hardware, operating system, or location. This tool also supports Informix Dynamic Server 9.x.
- **DB2 Table Editor:** Quickly and easily accesses, updates, and deletes data across multiple DB2 database platforms. This tool also supports Informix Dynamic Server 9.x

See [Resources](#) on page 53 for more information on these tools.

Section 3. DB2 tools

Tools overview

The tools that are included with DB2 (herein referred to as the *DB2 tools*; don't confuse them with the add-on DB2 tools, which are available separately) provide a whole array of time-saving, error-reducing graphical interfaces into most of the features of DB2. With these tools, you can perform the same tasks from a graphical user interface (GUI) that you can perform from a command line or API. With the tools, however, you do not have to remember complex statements or commands, and you get additional assistance through online help and wizards -- so let's hear it for the DB2 tools!

The DB2 tools are part of the *DB2 Administration Client*. When you install a DB2 server, you are actually installing all of the components of a DB2 Administration Client as well (though most people don't realize it). DB2 Administration Clients enable you to install DB2 tools on any workstation, and they allow you to manage remote database servers.

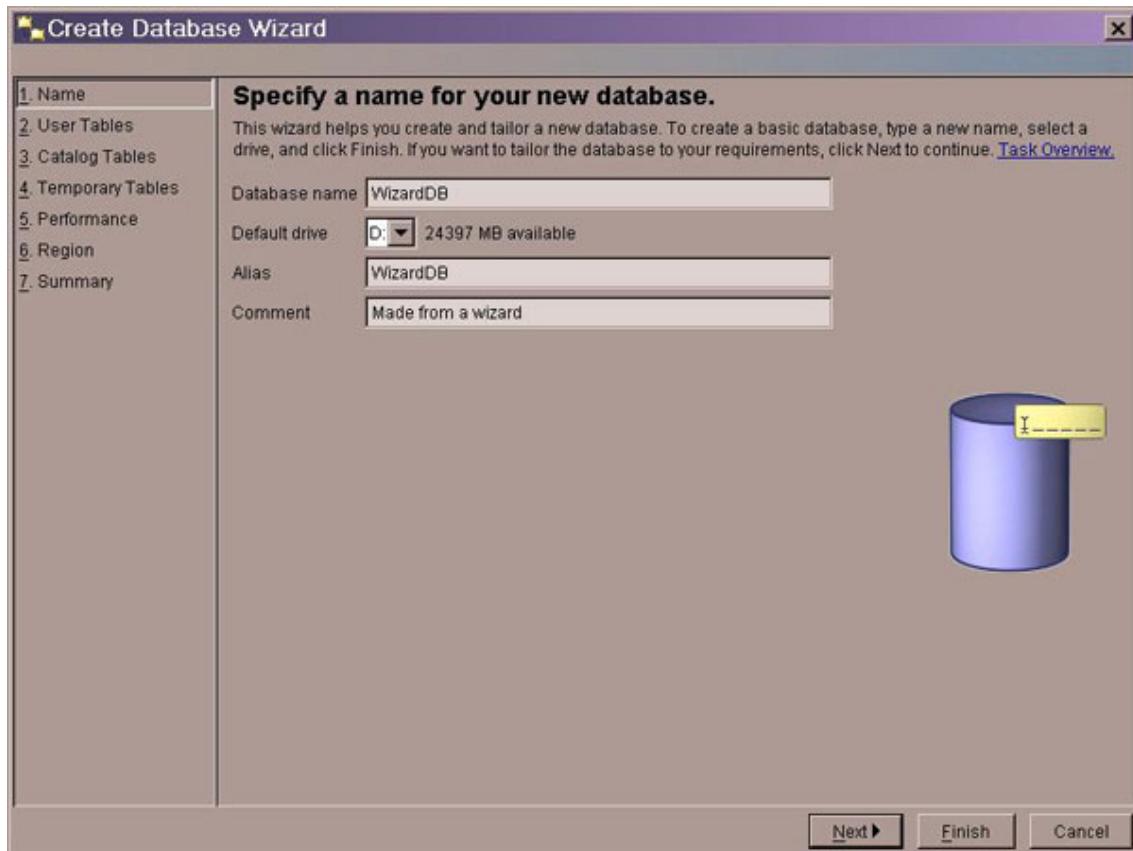
The DB2 tools are really divided into two camps: the *Control Center* (CC) and the *Configuration Assistant* (CA). The Control Center is primarily used for administering DB2 servers, and the Configuration Assistant is used for setting up client/server communications and maintaining registry variables (though it can do more; you will learn more about the CA in a bit). In addition, there are several other centers that are integrated and can be started from the Control Center.

Basic tool functions

There are about six basic features that you should be able to find in any DB2 tool (when applicable): *Wizards*, *Generate DDL*, *Show SQL/Show Command*, *Show Related*, *Filter*, and *Help*. Let's look at each in turn.

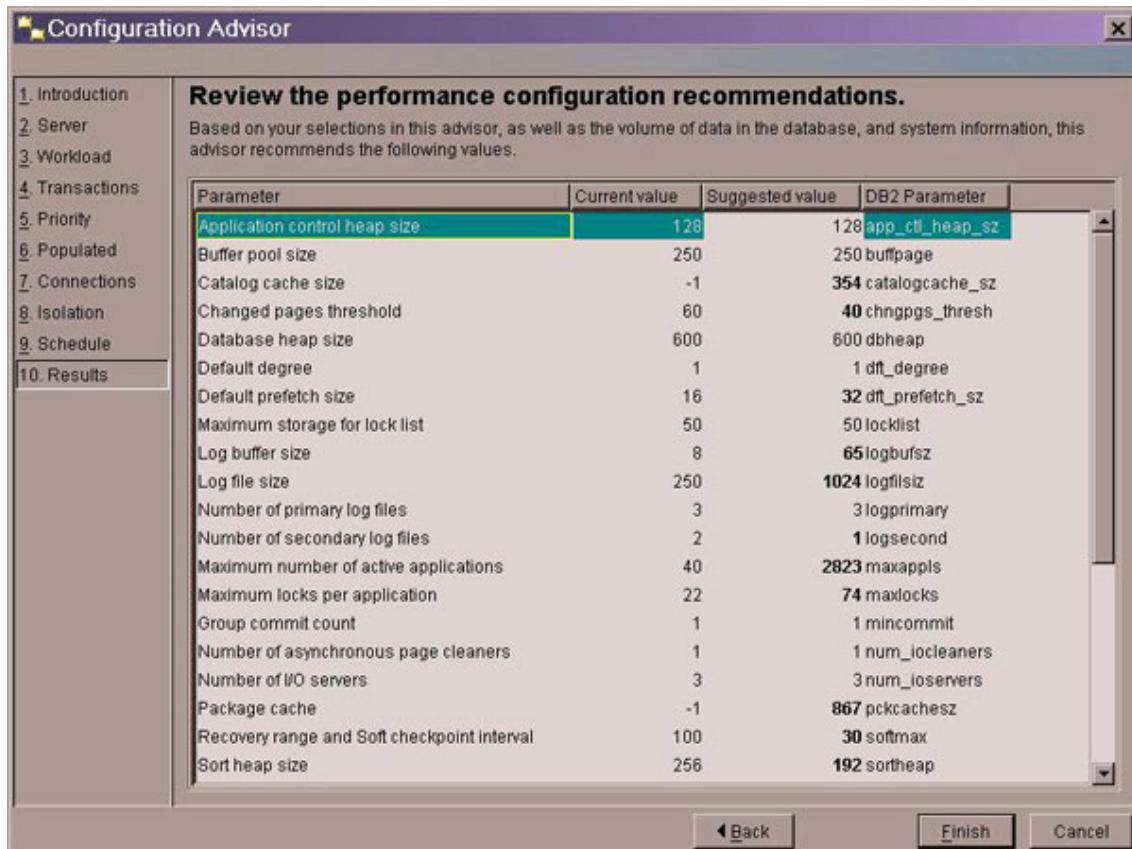
Wizards: Wizards can be very useful to both novice and expert DB2 users. Wizards help you complete specific administration tasks by taking you through each task one step at a time, recommending settings where applicable. Wizards are available through both the Control Center and the Configuration Assistant. They launch automatically when you perform specific actions; you can also launch them manually via the shortcut menu (right mouse click).

There are wizards for adding a database to your system (cataloging it), creating a database, backing up and restoring a database, creating tables, creating tablespaces, configuring two-phase commits, configuring database logging, updating your documentation, and tuning your performance. The following figure shows the Create Database wizard.

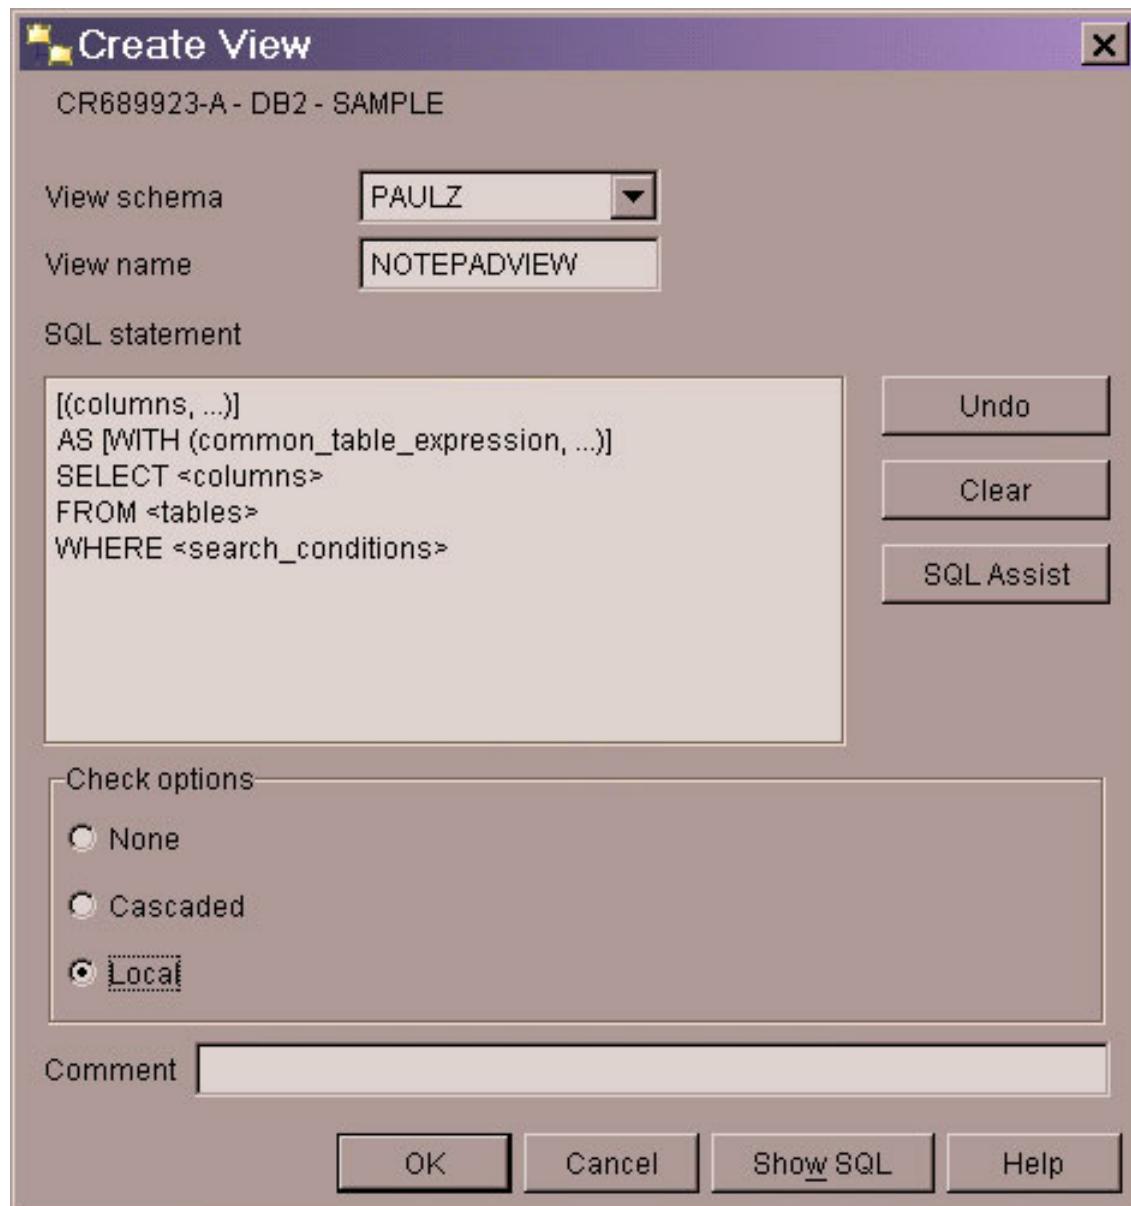


There are special kinds of wizards that are new to DB2 Version 8, called *advisors*. Traditional wizards take you step-by-step through a task, simplifying the experience by asking important questions or generating the complex command syntax for the action you want to perform. Advisors assist you with more complex activities, such as tuning and performance tasks, by gathering information and recommending options that you may not have considered. You can then accept or reject the advisor's advice. You can call advisors from the GUI as well as from APIs and the command-line interface.

Advisors are part of the IBM autonomic computing effort, which aims to make software and hardware more SMART (self-managing and resource tuning)! There are two advisors: the Configuration Advisor and the Design Advisor. For example, you can start the Configuration Advisor and respond to a number of questions. DB2 will then SMARTly consider the answers you gave and make a number of configuration recommendations based on your responses. The Design Advisor will consider a workload and suggest indexes that could help the performance for that workload. These functions reveal the difference between an advisor or a wizard: A wizard would help you create an index, but an advisor would actually suggest a specific index to create. Advisors truly let junior DBAs improve their skills and thereby reduce the effort and total cost of ownership of a DB2 solution. The following figure shows the suggested configuration parameter settings that the Configuration Advisor would suggest based on a particular set of responses.



Another type of wizard is a *notebook*. A notebook differs from the wizards you may be used to, since it doesn't step you through a particular process (e.g., creating a table); but notebooks do simplify tasks by eliminating the need to memorize clunky syntax. Notebooks exist for such tasks as setting up an event monitor or creating an index, buffer pool, trigger, alias, schema, or view. The following figure shows the Create View notebook.



When you take the exam, you should know about all the wizards, advisors, and notebooks, and how to use them. We recommend that you go through the Control Center and the Configuration Assistant, exploring these helpers and performing the various tasks with their help. Remember, practice makes perfect!

Generate DDL: The Generate DDL function allows you to re-create, and optionally save in a script file, the Data Definition Language (DDL), SQL statements, and statistics of database objects, authorization statements, tablespaces, nodegroups, buffer pools, database statistics, and pretty much anything else that makes up the basis of your database (except the data).

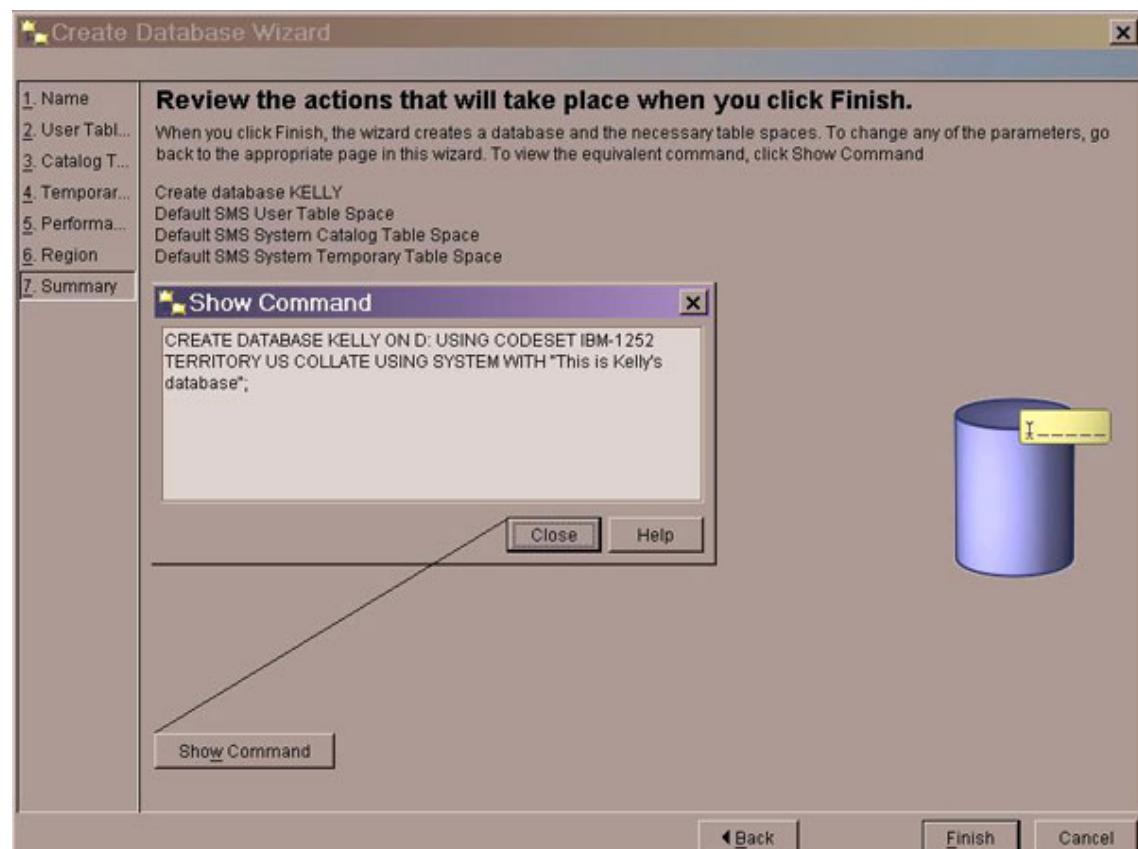
By using the Generate DDL feature, you can save the DDL to create identically defined tables, databases, and indexes in another database -- using it as a cookie cutter, if you will. Administrators like to use this option to create a test environment that mimics the production environment. A key feature of the Generate DDL option is its ability to create identical statistics for a test database *without* having to load the data in the tables. When you click on the Generate DDL option, you are actually running the

db2look DB2 system command.

If you want to move data into your new database objects, you could use the traditional LOAD or IMPORT utilities, or the db2move command. This tool facilitates the movement of large numbers of tables between DB2 databases located on distributed workstations. db2move queries the system catalog tables for a particular database and compiles a list of all user tables. It then exports these tables in PC/IXF format.

Show SQL/Show Command: If a tool generates SQL statements or DB2 commands, then the Show SQL or Show Command button will be available on that tool's interface. Clicking this button will show the actual statement or command that DB2 will use to perform the task you've requested. You can save the information returned by this function as a script for future reuse, or you can just use it to get a better idea of what is happening behind the interface (for all you UNIX folks who care about that sort of stuff).

You can save the output generated by the Show SQL or Show Command so you can schedule the commands as a script or create similar scripts without having to retype the statements or commands. You can also use the copy and paste features of your operating system to work with the generated syntax in another application. The following figure shows the CREATE DATABASE command that will be run based on the information entered in the Create Database wizard. (Of course, if the wizard were generating SQL, the option would be to show the SQL generated for the task.)



Show Related: The Show Related feature returns the immediate relationship between tables, indexes, views, aliases, triggers, tablespaces, user-defined functions (UDFs), and user-defined types (UDTs). For example, if you select a table and you choose to show the related views, you will only see the views that are based directly on that

specific base table. You will not see views that are based on the related views, because those views were not created directly from the table.

By seeing a list of related objects, you can better understand the structure of a database, determine what objects already exist in a database and their relationships to one another, and much more. For example, if you want to drop a table with dependent views, the Show Related feature will identify which views will become inoperative as a result of dropping the object.

The following figure shows the results of using the Show Related feature on a view. As you can see, the PAULZ.STAFFDEPT view has a dependency on the PAULZ.STAFF table. Using this information, you should be able to tell that, if the PAULZ.STAFF table were dropped, the PAULZ.STAFFDEPT view would become inoperable.

The screenshot shows a window titled "Show Related - STAFFDEPT". The title bar includes icons for minimize, maximize, and close. The main area displays the following information:

View - CR689923-A - DB2 - SAMPLE - PAULZ.STAFFDEPT

WARNING: Only first level dependencies are shown.

Below this, there is a toolbar with tabs: Aliases, Tables, Views, Triggers, and UDFs. The "Tables" tab is selected. A table header row follows:

Name	Schema	Type	Table space	Index table space	Long
------	--------	------	-------------	-------------------	------

One row of data is visible:

STAFF	PAULZ	T	USERSPACE1		
-------	-------	---	------------	--	--

At the bottom of the window are three buttons: Close, Show SQL, and Help.

Filter: You can filter the information that is displayed in the contents pane of any DB2 tool. You can also filter information that is returned from a query (such as limiting the number of rows in a result set).

The content pane's filter has been improved in DB2 Version 8. You can now save filters that you define, and recall them at a later time. If you click on the View button on the bottom right-hand corner of the Control Center, you will be presented with a pop-up dialog with which you can create, save, and edit filters. Take a moment now to create a filter for all of the database objects that you create under your own user ID. In later sections of this tutorial, you can then use this filter to quickly and easily find the database objects that you want to work with.

Help: I'm sure this doesn't need to be said, but the Help tool gives you help! Extensive help information is provided with the DB2 tools. A Help button exists on most dialog boxes as well as on the menu toolbar. You can get general help as well as help on how to fill out the fields and perform tasks of a particular tool. From the help menus, you can also access an index of terms used in the dialog or reference information, along with the information provided in the product manuals.

DB2 Version 8 features much improved documentation. Most noticeably, the information is now task oriented. Books are still provided in Adobe Acrobat's PDF format; however, the Information Center presents task-based help. In addition, there is a Documentation Update wizard that will update your existing DB2 information with each new FixPak. (This beats reading a 300-page README file four times a year!)

The DB2 processors: An introduction

The DB2 Command Line Processor (DB2 CLP), common to all DB2 products, is an application you can use to run DB2 commands, operating system commands, or SQL statements. This tool can be a somewhat cryptic method of invoking DB2 commands. However, the DB2 CLP can be a powerful tool because it extends its capability to store often-used sequences of commands or statements in batch files that can be run when necessary.

Some implementations of DB2 can use the operating system's native command-line interface to enter DB2 commands; others cannot. For this reason, we'll refer to two different processors in DB2: the DB2 Command Line Processor (DB2 CLP) and the DB2 Command Window (DB2 CW). You are likely to refer to them by the same name, however, and with good reason: they have the same icon in the Windows Start Menu and are both called DB2 CLP in the kernel in which they run. For this reason, users will often refer to the mode in which they do not have to prefix commands with the keyword db2 as the DB2 CLP in *interactive* mode.

In a Windows installation, both the DB2 CLP and the DB2 CW are located in the Command Line Tools folder, under the IBM DB2 folder.

The DB2 CLP allows you to enter DB2 commands interactively, without using the db2 prefix to tell the operating system that you're planning to enter a DB2 command. However, if you want to enter an OS command, you have to prefix it with the escape key (!). For example, in the DB2 CLP, if you wanted to run the dir command, you would enter !dir.

For all operating systems other than Windows, the DB2 CW is built into the operating

system's native CLP. In a Windows environment, you have to start a DB2 CW from a Windows command prompt by entering the db2cmd command or by selecting the appropriate option from the Start menu.

You can start the DB2 CLP from a DB2 CW by entering the db2 command on its own.

The following figure shows a command entered through the DB2 CW:

```
c:\>db2 connect to sample
Database Connection Information
Database server      = DB2/NT 8.1.0
SQL authorization ID = PAULZ
Local database alias  = SAMPLE

c:\>db2 select * from staff where dept=20
+----+-----+-----+----+-----+-----+
| ID | NAME  | DEPT | JOB | YEARS | SALARY | COMM |
+----+-----+-----+----+-----+-----+
| 10 | Sanders | 20 | Mgr |    7 | 18357.50 | - |
| 20 | Pernal  | 20 | Sales |   8 | 18171.25 | 612.45 |
| 80 | James   | 20 | Clerk |   - | 13504.60 | 128.20 |
| 190 | Sneider | 20 | Clerk |   8 | 14252.75 | 126.50 |
+----+-----+-----+----+-----+-----+
4 record(s) selected.

c:\>
```

Notice that I had to type in the keyword db2 to get this DB2 command to run. If I hadn't, the operating system would have thought that this was an operating system command, and would return an error. If you are using the DB2 CLP, you don't need to do this, as you can see in the figure below:

```
db2 => connect to sample
Database Connection Information
Database server      = DB2/NT 8.1.0
SQL authorization ID = PAULZ
Local database alias  = SAMPLE

db2 => select * from staff where dept = 20
+----+-----+-----+----+-----+-----+
| ID | NAME  | DEPT | JOB | YEARS | SALARY | COMM |
+----+-----+-----+----+-----+-----+
| 10 | Sanders | 20 | Mgr |    7 | 18357.50 | - |
| 20 | Pernal  | 20 | Sales |   8 | 18171.25 | 612.45 |
| 80 | James   | 20 | Clerk |   - | 13504.60 | 128.20 |
| 190 | Sneider | 20 | Clerk |   8 | 14252.75 | 126.50 |
+----+-----+-----+----+-----+-----+
4 record(s) selected.

db2 =>
```

Using the DB2 processors

When using a DB2 processor, you can use command-line options that alter the way the processor behaves. You can specify one or more processor options when you invoke a DB2 command. Some of the options that you can control are:

- Auto-commit each statement, defined by the `a` flag.
- Use an input file, defined by the `f` flag.
- Change the end-of-statement termination character (the default character is `:`), defined by the `t` flag.

You can get a list of all the valid options by entering the `list command options` in a DB2 processor.

There are two ways to change the options for a DB2 processor. You can set command options for a session by setting the `DB2OPTIONS` registry variable (which must be in uppercase), or by specifying command-line flags when you input a DB2 command. The latter method will override any settings made at the registry level.

To turn an option on, prefix the corresponding option letter with a minus sign (`-`); for example, to turn the auto-commit feature on (which is the default), enter `db2 -c create....`. To turn an option off, either surround the option letter with minus signs (`-c-`) or prefix it with a plus sign (`+`). Read the last two sentences again, because this can get confusing: a minus sign *before* a flag turns an option *on*, but a minus sign *before and after* a flag, or a *plus sign* before the flag, turns that option *off*. No, that isn't very intuitive (hey, I didn't write the code). Since this can be confusing, let's walk through an example with the auto-commit option.

By default, the auto-commit feature is set to (`-c`). This option specifies whether each statement is to be treated independently. When it is turned on, each statement is automatically committed or rolled back. If a statement is successful, it and all successful statements that were issued before it with auto-commit set to off (`+c` or `-c-`) are committed. If, however, the statement fails, it and all successful statements that were issued before it with autocommit set to off are rolled back. If auto-commit is set to off for the statement, you must explicitly issue a commit or rollback command.

In the following figure, we're changing the value of the auto-commit feature on the command line:

```
c:\>db2 create database test
DB20000I  The CREATE DATABASE command completed successfully.

c:\>db2 connect to test
Database Connection Information
Database server      = DB2/NT 8.1.0
SQL authorization ID = PAULZ
Local database alias  = TEST

c:\>db2 -c- create table a (c1 int)
DB20000I  The SQL command completed successfully.

c:\>db2 select c2 from a
SQL0206N "C2" is not valid in the context where it is used.  SQLSTATE=42703

c:\>db2 list tables
Table/View          Schema        Type Creation time
-----              -----        ---  -----
0 record(s) selected.

c:\>db2 connect reset
DB20000I  The SQL command completed successfully.

c:\>db2 drop database test
DB20000I  The DROP DATABASE command completed successfully.

c:\>_
```

So what happened? Well, first we created a database named TEST. After connecting to it, we created a table called A, but we did this while at the same time turning the default auto-commit option to off by using the `-c-` option. (We could have prefixed this flag with a plus sign `+c` and it would have done the same thing.) After creating table A (but not committing this action, remember), we selected values from column C2. We received an error because there is no column C2 in table A. This error causes a rollback of all non-committed work, so the table creation statement that we ran is rolled back, and now table A no longer exists (since it was created with the auto-commit feature turned off). You can see that I verified this with the DB2 LIST TABLES command.

Try the exact same sequence of commands, only this time use the `+c` option. You should experience the same behavior. After that, try it with the `-c` option (which would be the same as not specifying any option at all, since auto-commit is the default). This time the table should still exist after the error, since the auto-commit feature is set to on and that means that each and every statement either commits or rolls back.

You can find more detailed information about command-line processor options in the *DB2 Command Reference*.

Your operating system may have a maximum number of characters that it can read in any one statement (even when it wraps to the next line in your display). To work around this limitation when entering a long statement, you can use the line continuation character (`\`). When DB2 encounters the line continuation character, it reads the next line and concatenates the two lines during processing. You can use this character with both DB2 processors. The figure below illustrates its use in the DB2 CLP:

Using the line continuation character (\) for long statements - db2

```
db2 => connect to sample
Database Connection Information
Database server      = DB2/NT 8.1.0
SQL authorization ID = PAULZ
Local database alias = SAMPLE

db2 => select * from employee \
db2 <(cont.) => where job='MANAGER'

EMPNO FIRSTNAME MIDINIT LASTNAME      WORKDEPT PHONENO HIREDATE   JOB
EDLEVEL SEX BIRTHDATE SALARY    BONUS      COMM

000020 MICHAEL     L    THOMPSON      B01      3476    10/10/1973 MANAGER
18 M 02/02/1948   41250.00    800.00   3300.00
000030 SALLY       A    KWAN          C01      4738    04/05/1975 MANAGER
20 F 05/11/1941   38250.00    800.00   3060.00
000050 JOHN        B    GEYER          E01      6789    08/17/1949 MANAGER
16 M 09/15/1925   40175.00    800.00   3214.00
000060 IRVING      F    STERN          D11      6423    09/14/1973 MANAGER
16 M 07/07/1945   32250.00    500.00   2580.00
000070 EVA         D    PULASKI      D21      7831    09/30/1980 MANAGER
16 F 05/26/1953   36170.00    700.00   2893.00
000090 EILEEN      W    HENDERSON    E11      5498    08/15/1970 MANAGER
16 F 05/15/1941   29750.00    600.00   2380.00
000100 THEODORE    Q    SPENSER      E21      0972    06/19/1980 MANAGER
14 M 12/18/1956   26150.00    500.00   2092.00

? record(s) selected.

db2 =>
```

When you are using DB2 CW to enter commands, you may run into a problem with some of the special characters shown below:

\$ & * () ; < > ? \ ' "

The operating system shell may misinterpret these characters. (Of course, this is not an issue in the DB2 CLP, since it is a separate application specifically designed for DB2 commands.) You can circumvent system operators that you want to be interpreted by DB2 and not by the operating system by placing your entire command within quotes, as follows:

```
db2 "select * from staff where dept > 10"
```

Try entering the previous command in a DB2 CW without the quotes. What happened? Look at the contents of the directory where you issued the command. I'll bet you'll find a file there called 10 that has an SQL error in it. Why? Well, DB2 interpreted your SQL command `select * from staff where dept` and tried to place those contents in a file called 10. The `>` sign piped any output to the file 10 as per the native operating system syntax. `select * from staff where dept` is of course an incomplete SQL statement; hence the error. The incorrect results are due to the operating system misinterpreting the special character.

You should experiment with both DB2 processors to get a feel for which is the best choice in any given circumstance. Take some time now to enter commands in both the DB2 CLP and the DB2 CW.

Section 4. The Control Center

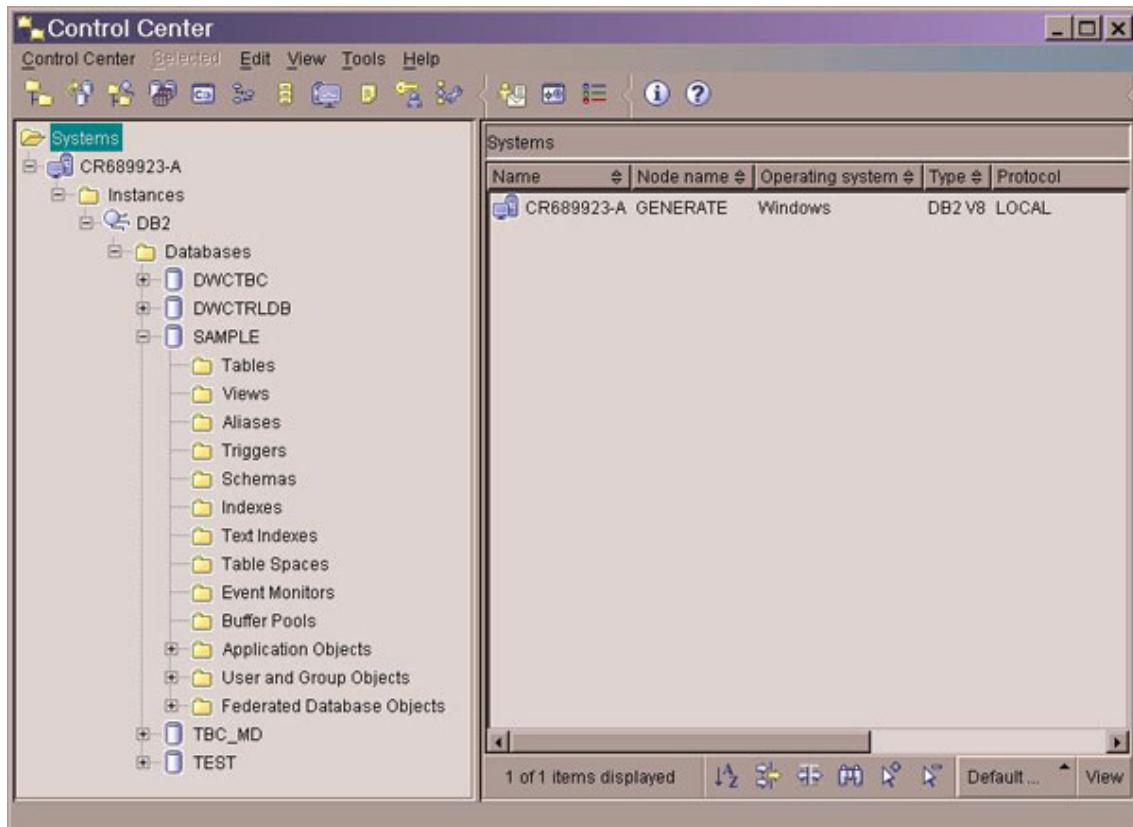
Overview

The Control Center (CC) is the central management tool for DB2 servers. You can use the CC to manage systems, DB2 instances, DB2 z/OS subsystems, IMS data sources, databases, database objects, and much, much more. From the CC, you can also open other centers and tools to help you optimize queries, schedule jobs, and write and save scripts; perform data warehousing tasks; create stored procedures and user-defined functions; work with DB2 commands; and monitor the health of your DB2 system. To start the CC on Windows, in the Start menu, select **Control Center** from the General Administration Tools folder in the IBM DB2 folder. You can also enter the db2cc command from a command prompt on any operating system to start the CC.

Among other tasks, a database administrator can use the Control Center to:

- Add DB2 systems, federated systems, DB2 for z/OS subsystems, IMS systems, and local and remote instances and databases to the object tree for management.
- Manage database objects. You can create, alter, and drop databases, tablespaces, tables, views, indexes, triggers, and schemas. You can also manage users.
- Manage data. You can load, import, or export data, reorganize data, and collect statistics.
- Perform preventive maintenance by backing up and restoring databases or tablespaces.
- Schedule jobs to run unattended. To schedule tasks through the CC, you must first create a TOOLS catalog database. If you did not create this database when you installed DB2 (it is an optional part of the installation process), you can do so from the Tools Settings option in the Tools action menu bar. If you haven't done so already, create it now.
- Configure and tune instances and databases.
- Manage database connections.
- Monitor and tune performance. You can run statistics, look at the execution path of a query, start event and snapshot monitoring, generate SQL or DDL of database objects or commands, and view relationships between DB2 objects.
- Troubleshoot.
- Manage data replication.
- Manage applications.
- Manage the health of your DB2 system.
- Launch other DB2 centers.

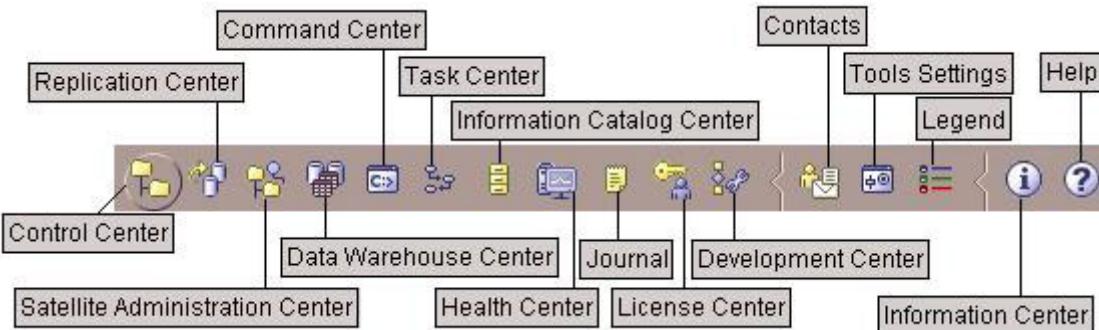
The Control Center is shown below. As you can see, if you have experience with any Microsoft tools, you can easily use this tool; the layout, with objects on the left and details on the right, should be familiar to you.



The Control Center relies on the Database Administration Server (DAS). DAS helps the Control Center schedule jobs to run against a database server, manage objects on remote database servers, and more.

To see everything that you can do with the Control Center, right-click on any object from the object tree. A pop-up menu shows all the functions you can perform on the selected object. For example, on the Tables folder, you can create a new table, create a filter for what is displayed in the contents pane, or refresh the view. The tasks that you can perform depend on the object that you select. I strongly recommend that you go through each folder and object and right-click your way to familiarity!

From the launchpad, you can access all the other DB2 centers that are integrated into the Control Center. If you select the Control Center from the launchpad, another instance of this tool is started. If you install the DB2 Everyplace SyncServer, then the Mobile Devices Administration Center (MDAC) will be added to the launchpad. The launchpad for a typical DB2 installation is shown below:

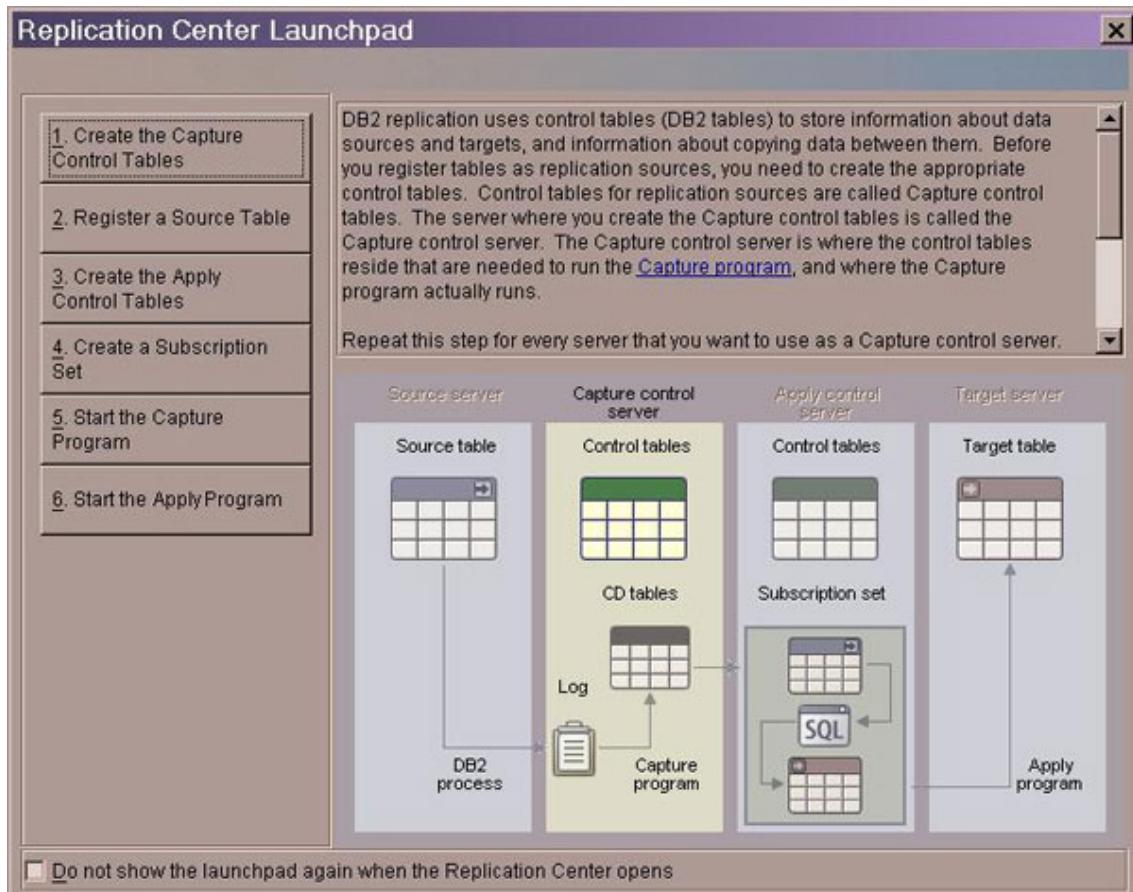


The next several panels in this section present a detailed description of the Control Center tools that can be started from the launchpad. Each tool that you launch includes the launchpad, so you can launch any DB2 tool from any other DB2 tool. This section covers the DB2 tools and centers that are used most often. The following sections in this tutorial will cover the remaining tools (for example, the Data Warehouse Center) available from the launchpad that are not covered in this section.

The DB2 Replication Center

Use the DB2 Replication Center (DB2 RC) to administer replication between a DB2 database and other relational databases (DB2 or non-DB2). From the DB2 RC, you can define replication environments, apply designated changes from one location to another, and synchronize the data in both locations.

You can start the DB2 RC from the Start menu, from a DB2 tool's launchpad, or by entering the `db2rc` command at the command prompt. The following figure gives you an idea of what the Replication Center looks like:



A replication-specific launchpad (depicted in the figure above) is available to guide you through some of the basic replication functions. Some of the key tasks that you can perform with the Replication Center include:

- Create replication control tables
- Register replication sources
- Create subscription sets
- Operate the Capture program
- Operate the Apply program
- Monitor the replication process
- Perform basic troubleshooting for replication

The DB2 Satellite Administration Center

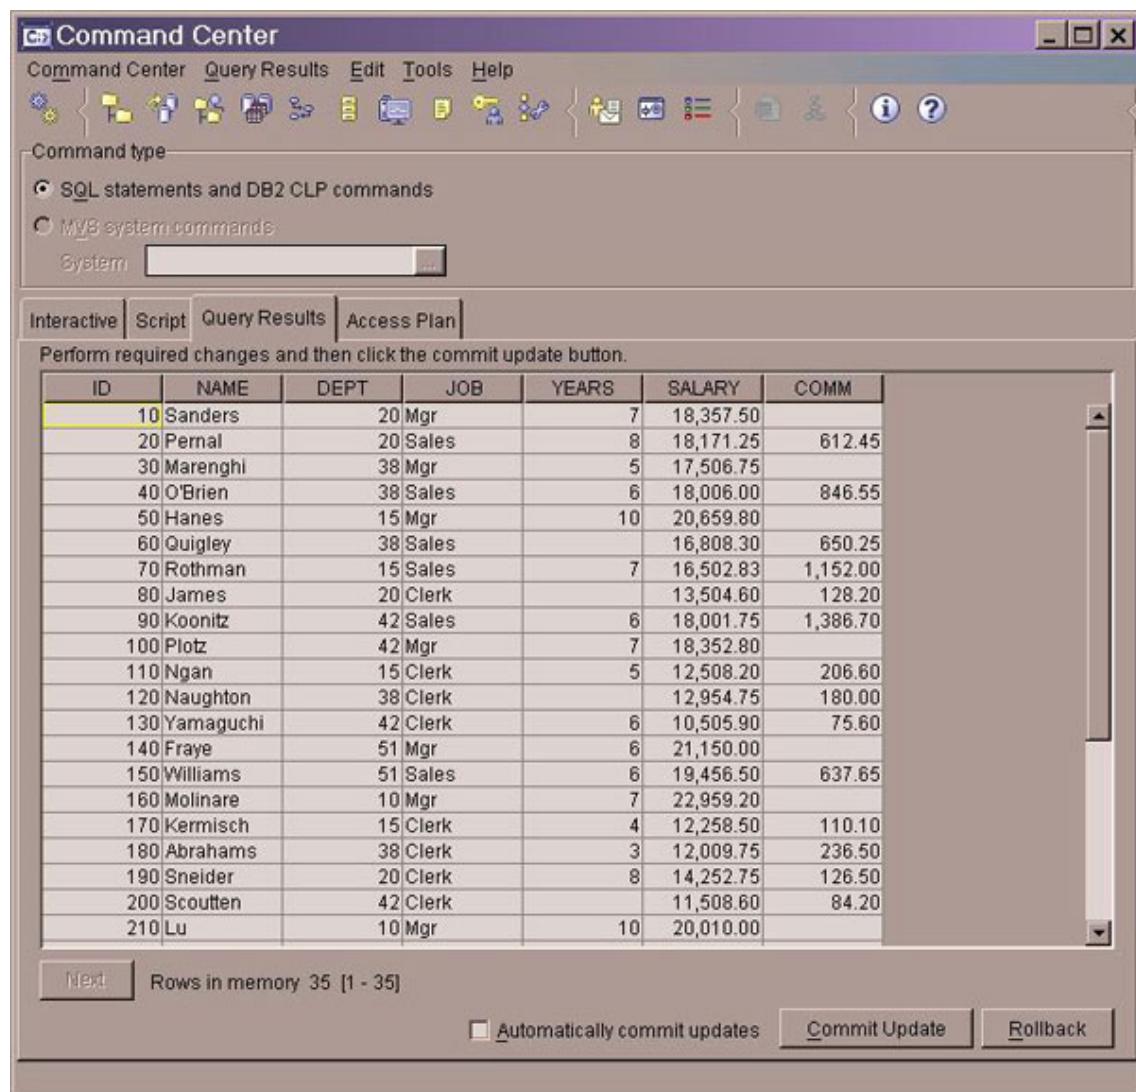
Use the Satellite Administration Center (DB2 SAC) to set up and administer a group of DB2 servers that perform the same business function. These servers, known as *satellites*, all run the same application and have the same DB2 configuration (database definition) to support that application. You can start the DB2 SAC from the launchpad in any DB2 tool.

The DB2 Command Center

Use the DB2 Command Center to execute DB2 commands and SQL statements to; execute MVS console commands; work with command scripts; view a graphical representation of the access plan for an SQL statement.

You can start the Command Center from the Control Center by clicking the appropriate icon on the toolbar, or entering the `db2cctr` command at the command prompt. In a Windows environment, you can also start the Command Center from the Start menu.

Take a look at the DB2 Command Center in the figure below:



Different tabs on the DB2 Command Center provide different features:

- On the **Interactive tab**, you can execute SQL statements or DB2 commands. (Entering DB2 commands in the Command Center is like working in interactive DB2 CLP mode: you don't need to use the `db2` prefix.) To run a DB2 command or statement that you enter, click on the gears (located in the top right-hand corner of

this tool) or press Ctrl+Enter. You can also enter operating system-specific commands from the DB2 Command Center by preceding the command with a bang (!) sign. For example, to list the contents of the current directory, enter !dir.

- On the **Script tab**, you can execute commands in sequence, create and save scripts, store a saved script in the Task Center (where you can schedule the script to run at a specific time -- more on this on the next panel), run an existing script, or schedule a task.
- On the **Query Results tab**, you can see the results of your query. You can also save the query's results or edit the contents of the table. (This page is the one shown in the figure above.)
- On the **Access Plan tab**, you can see the access plan for any explainable statement that you specified on the Interactive or Script pages. DB2 generates the access plan when it compiles the SQL statement. You can use this information to tune your queries for better performance. If you specify more than one statement, an access plan is created only for the first statement.

The DB2 Command Center also comes with the SQL Assist tool (which I'll talk about more in the next section). To invoke the SQL Assist tool, click the SQL Assist button on the Interactive page.

The DB2 Command Center can also be run in a Web-based mode. In this mode, any Web browser, PDA, mobile, or other pervasive device that has access to the Internet can execute commands against a DB2 server. This helps DBAs continually keep in touch with their DB2 systems. In Web mode, the Command Center does not have the Visual Explain or SQL Assist features.

The DB2 Task Center

Use the DB2 Task Center (DB2 TC) to run tasks, either immediately or according to a schedule, and to notify people about the status of completed tasks. The DB2 TC is the replacement tool for the DB2 Version 7 Script Center, and includes all of that tool's functionality and more. You can start the DB2 TC from the Start menu in a Windows environment, from a DB2 tool's launchpad, or by entering the db2tc command from a command prompt.

A *task* is a script accompanied by associated failure or success conditions, schedules, and notifications. You can create a task within the DB2 TC, create a script within another tool and save it to the DB2 TC, import an existing script, or save the options from a DB2 dialog or wizard (such as the Load wizard) as a script. The script can contain DB2, SQL, or operating system commands.

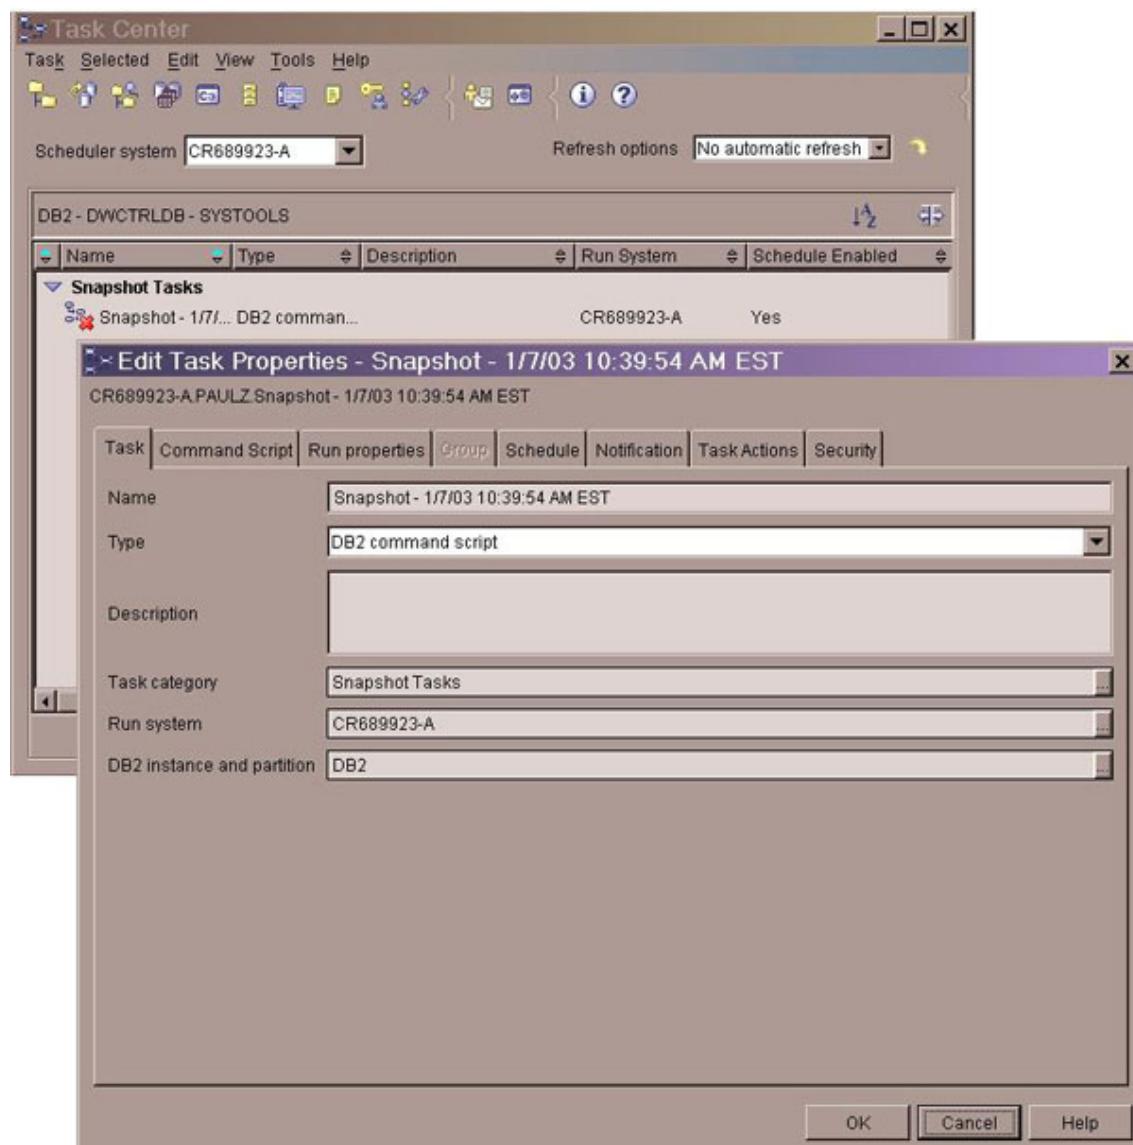
For each task, you can:

- Schedule the task
- Specify success and failure conditions
- Specify actions that should be performed when this task completes successfully or when it fails

- Specify e-mail addresses (including pagers) that should be notified when this task completes successfully or when it fails

You can also create a *grouping task*, which combines several tasks into a single logical unit of work. When a grouping task meets the success or failure conditions that you define, any follow-on tasks are run. For example, you could combine three backup scripts into a grouping task and then specify a reorganization as a follow-on task that will be executed if all of the backup scripts execute successfully. All of these features make the DB2 Task Center an indispensable resource for DBAs charged with managing a DB2 environment.

Take a look at the Task Center in the figure below:



The DB2 Health Center

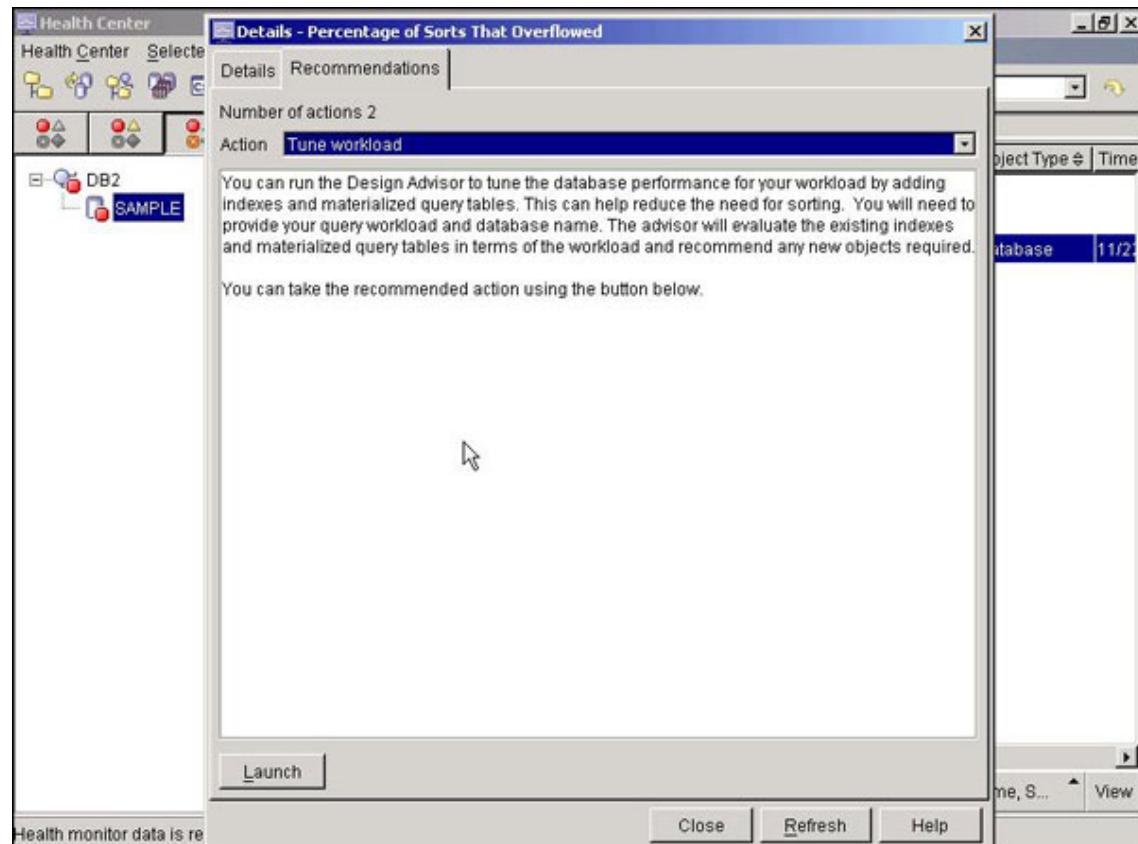
Use the DB2 Health Center (DB2 HC) to monitor the state of the DB2 environment and make any necessary changes to it. You can start the DB2 HC from the Start menu in a Windows environment, from any DB2 tool's launchpad, or by entering the db2hc command at a command prompt.

When you use DB2, a monitor continuously keeps track of a set of health indicators. If the current value of a health indicator is outside the acceptable operating range defined by its warning and alarm thresholds, the health monitor generates a health alert. DB2 comes with a set of predefined threshold values, which you can customize. For example, you can customize the alarm and warning thresholds for the amount of space used in a tablespace.

Depending on the configuration of the DB2 instance, some or all of the following actions can occur when the health monitor generates an alert:

- An entry is written in the administration notification log, which you can read from the Journal (see the next panel for details).
- The health center status beacon appears in the lower right corner of the DB2 GUI Tool's window.
- A script or task is executed.
- An e-mail or pager message is sent to the contacts that you specify for this instance.

The figure below depicts the DB2 HC and a suggested resolution strategy in response to a condition it has detected.



There are a number of key tasks that you can perform with the DB2 HC. For instance,

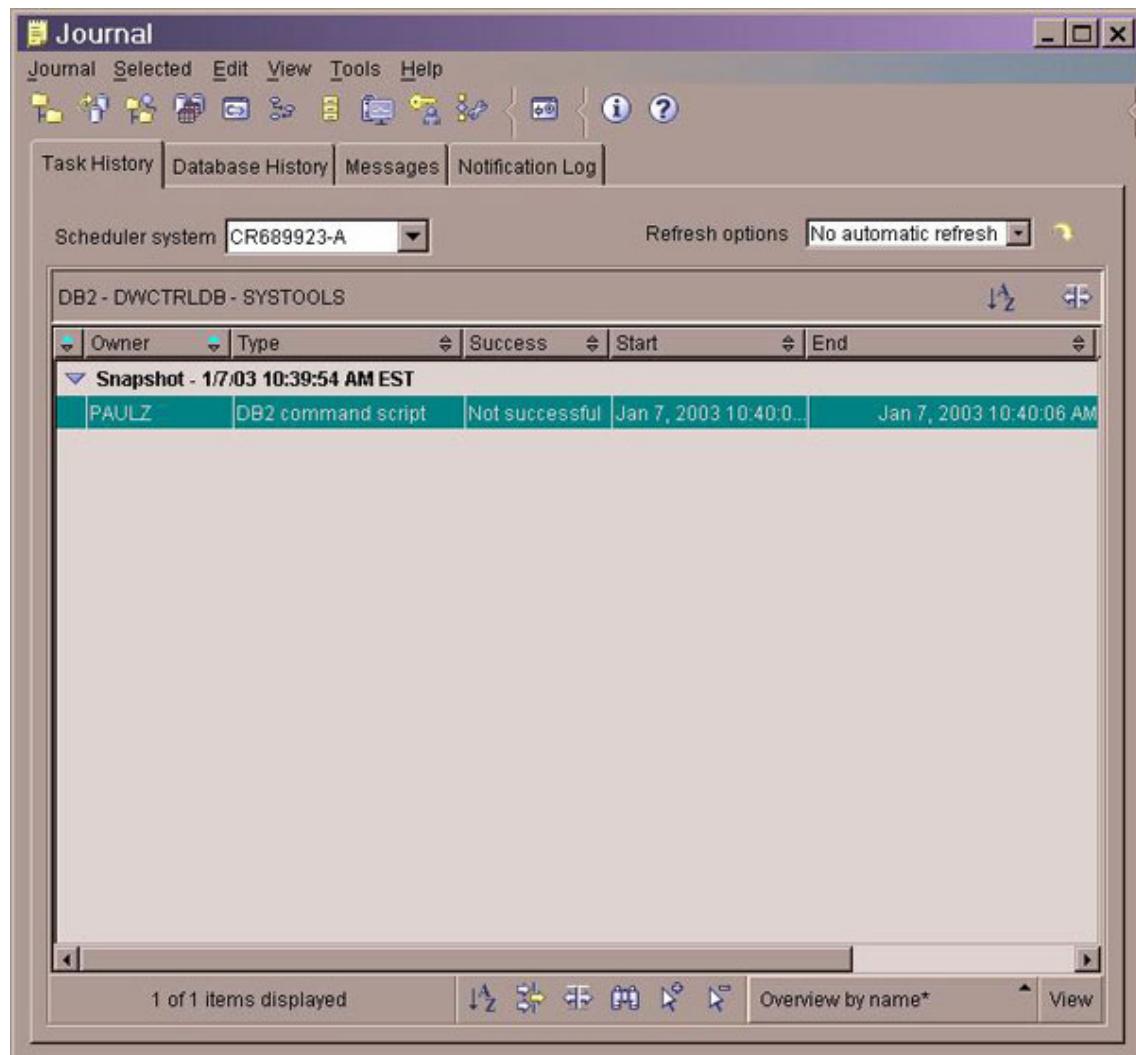
you can:

- View the status of the DB2 environment. Beside each object in the navigation tree, an icon indicates the status for that object (or for any objects contained by that object). For example, a green diamond icon beside an instance means that the instance and the databases contained in the instance do not have any alerts.
 - View alerts for an instance or database. When you select an object in the navigation tree, the alerts for that object are shown in the pane to the right.
 - View detailed information about an alert and recommended actions. When you double-click an alert, a notebook appears. The first page shows the details for the alert; the second shows any recommended actions. (This second page is shown in the figure above.)
 - Configure health monitor settings for a specific object, or the default settings for an object type or for all the objects within an instance.
 - Select which contacts will be notified of alerts via an e-mail or pager message.
 - Review the history of alerts for an instance.
-

The DB2 Journal

The DB2 Journal displays historical information about tasks, database actions and operations, Control Center actions, messages, and alerts. You can start the DB2 Journal from the Start menu in a Windows environment, or from the launchpad of any DB2 tool.

The figure below depicts the Journal, with some information from past events displayed:



There are four tabs in this tool, each providing the DBA with valuable information:

- **Task History:** Shows the results of tasks that were previously executed. You can use this information to estimate how long future tasks will run. This page contains one row for each execution of a task. In contrast, the Task Center contains only one row for each task no matter how many times the task is executed.

For each completed execution of a task, you can perform the following actions:

- View the execution results
- View the task that was executed
- Edit the task that was executed
- View the task execution statistics
- Remove the task execution object from the Journal

The Task History page is illustrated in the figure above.

- **Database History:** Shows information from the recovery history file. This file is updated when various operations are performed, including: backup, restore, roll forward, load, and reorganization. This information could be useful if you need to restore a database or tablespace.

- **Messages:** Shows messages that were previously issued from the Control Center and other GUI tools.
 - **Notification Log:** Shows information from the administration notification log.
-

The DB2 License Center

The DB2 License Center displays the status of your DB2 license and usage information for the DB2 products installed on your system. In addition, it enables you to configure your system for proper license monitoring. You can use the License Center to add new licenses, set a concurrent user policy, upgrade a try-and-buy license to a production license, and much more. You can also control DB2 licenses through the command line using the db2licm command.

Contact List

This tool is used to set up a list of contacts whom the system should alert (via e-mail or pager) if a threshold that you've defined in the Health Center has been exceeded, or is approaching a cautionable state. You can define when you want the system to send out a notification.

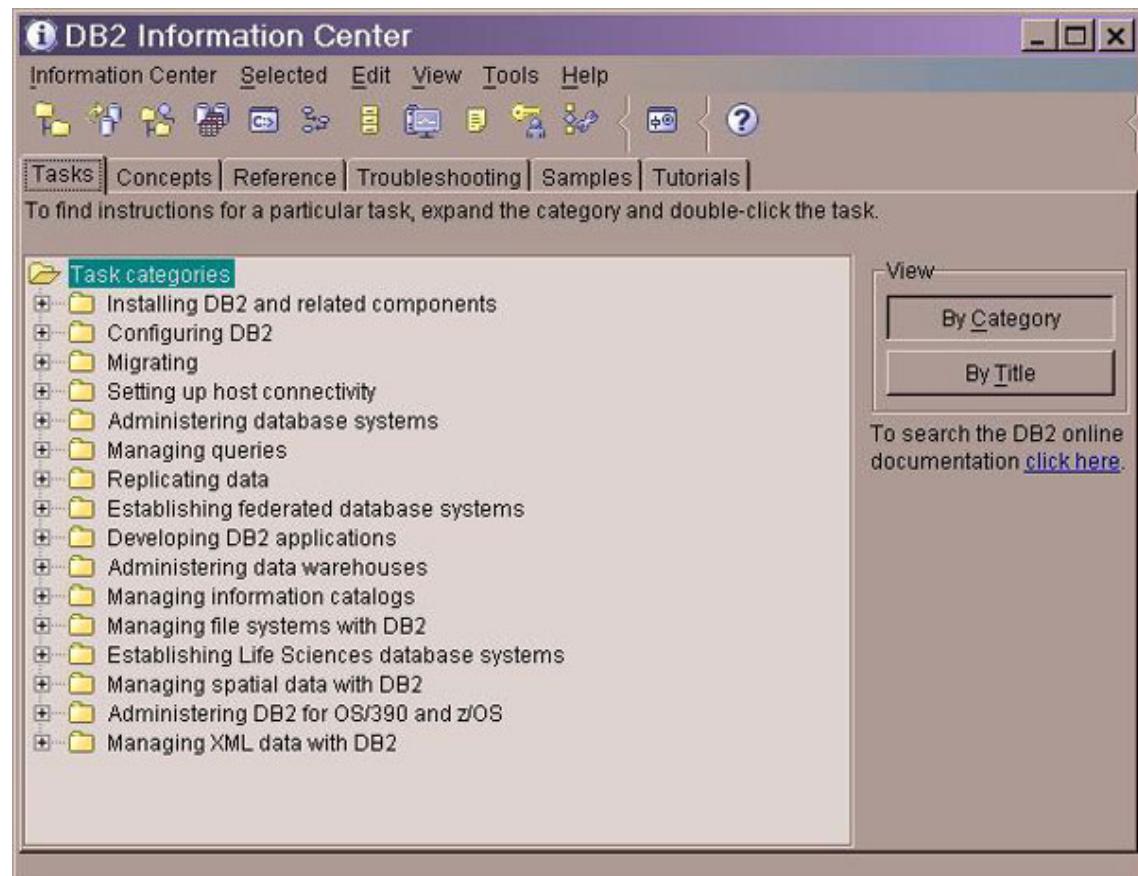
The DB2 Information Center

Use the Information Center to find information about tasks, reference material, troubleshooting, sample programs, and related Web sites.

This center will prove to be your one-stop shop for DB2 information. You can start the Information Center from the Control Center, from the Start menu in a Windows environment, or by entering the db2ic command.

The Information Center and the DB2 documentation have been completely reworked in Version 8 to be more task oriented. There is even a Documentation Update wizard that will update all of your existing documents with each FixPak, so you no longer need to read through each new README to learn about new features.

The Information Center provides six types of information, each shown as a tab in the figure below.



When you select an item in one of the lists, the Information Center launches a viewer to display the information. The viewer might be the system help viewer, an editor, or a Web browser, depending on the kind of information you select.

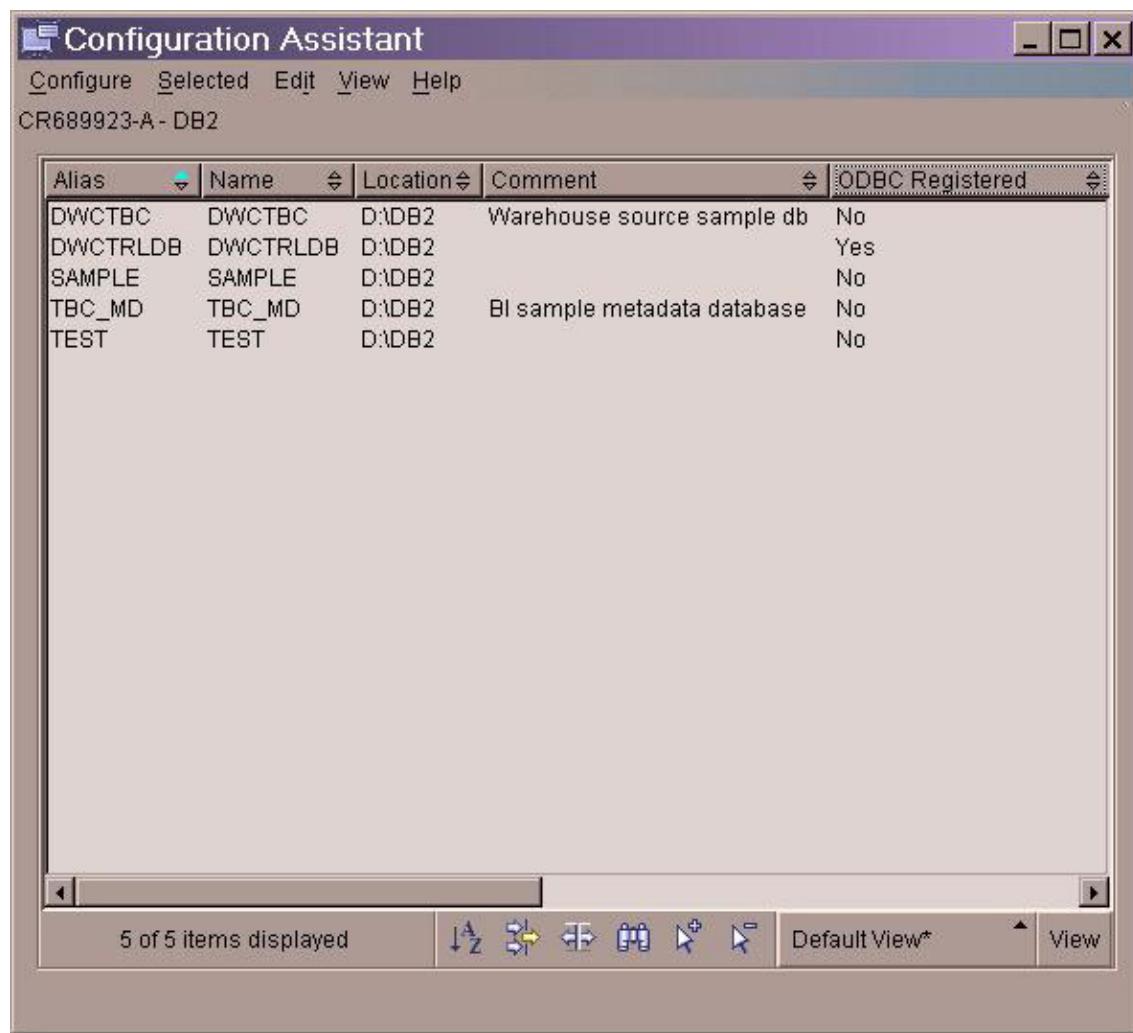
The Information Center also provides a new Java search tool to help you look for specific topics without browsing all of the information lists.

Section 5. The Configuration Assistant

Overview

The Configuration Assistant (CA) is a new tool in DB2 Version 8 that lets you maintain a list of databases to which your applications can connect, manage, and administer. The CA is the successor to the Client Configuration Assistant (CCA) available in DB2 Version 7.

The CA has a new interface and vast array of new functions; it is also now written completely in Java. In DB2 Version 7, the CCA only ran on Windows and OS/2 workstations (OS/2 is not longer a supported DB2 server platform in Version 8). Users who choose to leverage the scalability and robustness of UNIX or Linux can now use this fantastic tool. You can start the CA by entering the db2ca command on a command prompt, or, in a Windows environment, from the Start menu. Take a look at the CA in the figure below:



Each database that you wish to access from the CA must be cataloged at a DB2 client before you can work with it. Use the CA to configure and maintain database objects

that you or your applications will be using. The Add Database wizard will help you catalog nodes and databases, while shielding you from the inherent complexities of these tasks.

From the CA, you can work with existing databases, add new ones, bind applications, set client database manager configuration parameters, and import and export configuration profiles. The CA's graphical interface makes these complex tasks easier through:

- Wizards that help you perform certain tasks
- Dynamic fields that are activated based on your input choices
- Hints that help you make configuration decisions
- The Discovery feature, which can retrieve information that is known about databases that reside on your network

As you can see in the figure above, the CA displays a list of the databases to which your applications can connect and from which it can access data. Each database is identified first by its database alias, then by its name. You can use the Change Database wizard to alter the information associated with databases in this list. The CA also has an Advanced view, which uses a notebook to organize connection information by the following objects:

- Systems
- Instance nodes
- Databases
- Database Connection Services (DCS)
- Data sources

You can use the CA to catalog databases and data sources (like CLI and ODBC parameters), configure your instances, and import and export client profiles.

Take the time now to create a database and add it to your connection list. Try to use the discovery feature to add your new database for local access. Now go through each of the other features in the CA to understand the framework in which the CA presents them.

Section 6. Other DB2 tools

Other DB2 tools

There are a multitude of other DB2 tools you can leverage to make your work easier. These tools are delivered free of charge as standalone tools, in the Control Center, or the Configuration Assistant. Don't confuse these tools with the add-on DB2 tools that are available separately, which you learned about in Section 1 of this tutorial (see [DB2 add-on tools](#) on page 10).

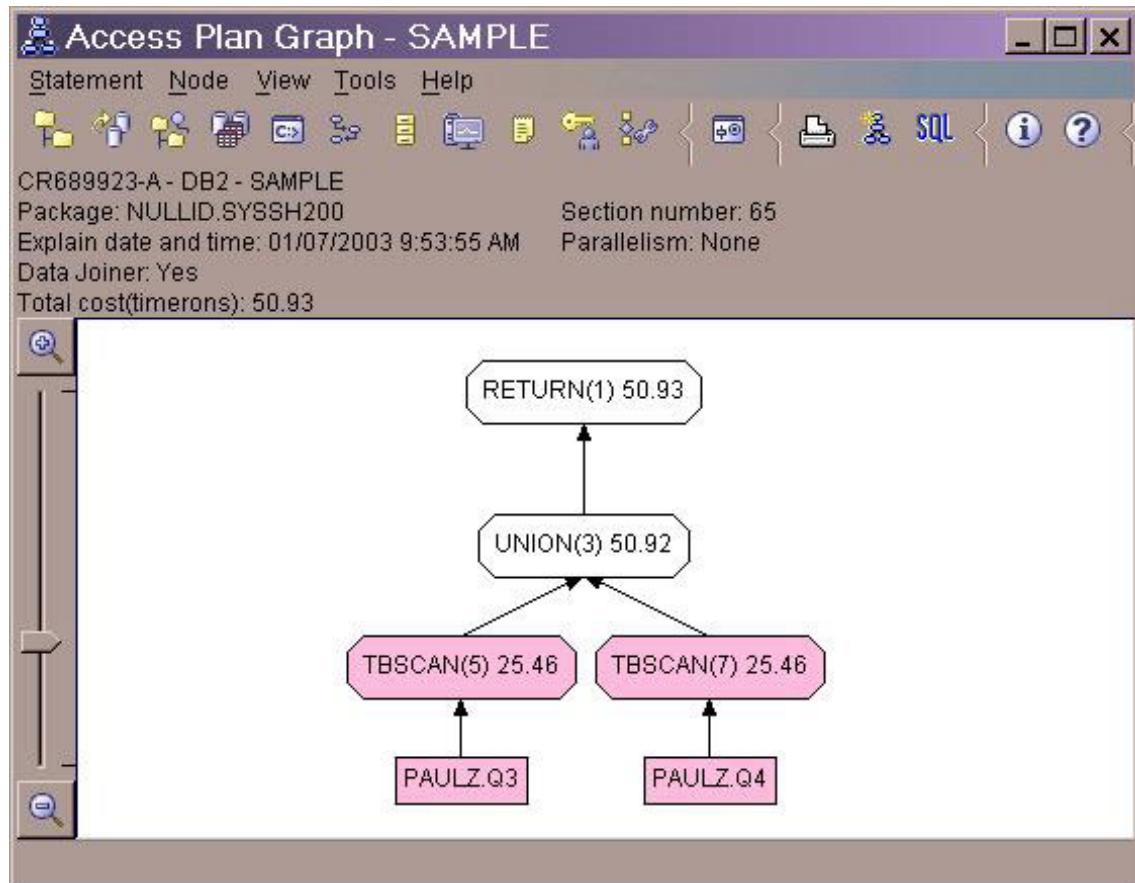
Visual Explain

Visual Explain lets you view the access plan for an explained SQL statement as a graph. You can use the information available from the graph to tune your SQL query for better performance. Visual Explain also lets you dynamically explain an SQL statement and view the resulting access plan graph.

The DB2 optimizer chooses an access plan and Visual Explain displays this plan, in which tables and indexes -- and operations on them -- are represented as nodes, and the flow of data is represented by the links between nodes.

The best part of Visual Explain is that you don't even have to run the query to get the information you are looking for. For example, let's say you suspect a query you have is inefficiently written; using Visual Explain, you can graphically look at the cost of the query without actually running it.

You can get the graphical access plan for a query without running it by entering it in the Control Center. From the Control Center tree view, select the database you want to work with, right-click, and select **Explain SQL**. Enter the SQL statement that you want to explain and click **OK**. An example of a visually explained query is shown in the figure below:



The Snapshot and Event Monitors

There are two utility monitors provided in DB2 to help you better understand your system and the impact of operations upon it.

The *Snapshot Monitor* captures database information at specific points in time. You determine the time interval between these points and the data that will be captured. The Snapshot Monitor can help analyze performance problems, tune SQL statements, and identify exception conditions based on limits or thresholds. In DB2 Version 8, you now have the ability to retrieve snapshot information into a DB2 table via an SQL UDF. In DB2 Version 7, the only way to get this type of information was programmatically using a C API (which is still supported in DB2 Version 8).

The *Event Monitor* is a tool that enables you to analyze resource usage by recording the state of the database at the time that specific events occur. For example, you can use the Event Monitor when you need to know how long a transaction has taken to complete, or what percentage of available CPU resources an SQL statement has used.

Tool Settings

The Tools Settings notebook allows you to customize the DB2 graphical tools and some of their options. You can use this notebook to:

- Set general property settings, such as the termination character, the automatic startup of DB2, hover help and infopop characteristics, the maximum number of rows in a result set, and more
 - Change the fonts for menus and text
 - Set DB2 for z/OS Control Center properties
 - Configure Health Center notification
 - Set up a default DB2 scheduling scheme for scheduled tasks
-

DB2 Development Center

The DB2 Development Center (DB2 DC) is the successor to the DB2 Version 7 Stored Procedure Builder (DB2 SPB). The DB2 DC builds upon the DB2 SPB with *lots* of new features, functions, and benefits that are bound to make developers smile.

In a nutshell, the DB2 DC is a rapid iterative development environment for building stored procedures (SPs), user-defined functions (UDFs), structured data types, and much more. This tool was completely rewritten to allow for concurrent task execution, flexible docking, enhanced scalability, and higher productivity. This Integrated Development Environment (IDE) can stand on its own, or be embedded into WebSphere Studio Application Developer or any of the Microsoft Visual Studio development products. As a result, developers can build DB2 business logic *without* leaving their favorite IDE -- a feature they'll love!

Some of the Development Center features include:

- **Functionality:**

- Support for the entire family of DB2 server operating systems, including z/OS, OS/400, UNIX, and Windows
- Enhanced z/OS support, including specialized SQL IDs (package owner, build owner, and secondary SQL ID) and advanced build options
- Support for developing SQL and Java stored procedures, SQL scalar and table user-defined functions, MQSeries, OLE DB, and XML table functions, and structured data types for EJB methods and properties
- Support for viewing live database tables, views, triggers, SPs, and UDFs

- **Usability:**

- A quick-start launchpad for guiding novice users through the initial set of development tasks
- A customizable workbench that uses flexible docking
- Separation of project-specific server objects, presented in the project view, from live database objects, presented in the server view

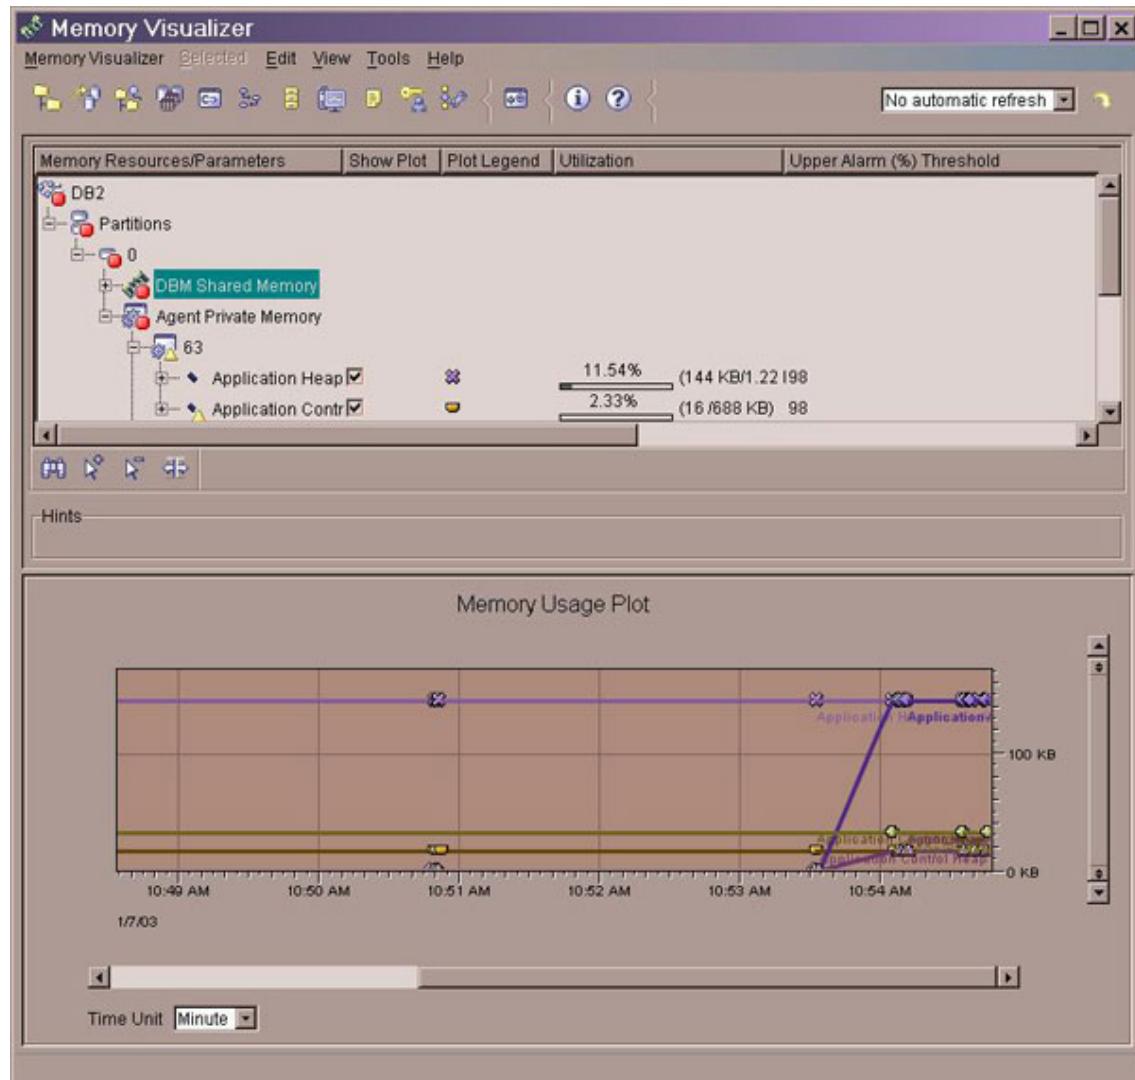
- **Project management:**

- Multiple concurrent project development
- Import of SPs and functions from databases, other projects, and other source files
- Export of project files to a deployment project or script
- A deployment wizard, tool, and command line, with binary deployment support for similar DB2 servers
- **Testing and debugging**
 - Ability to test SPs and UDFs written in any language
 - Enhanced round-trip debugging of SQL SPs with variable value change support using an integrated SQL debugger
 - Simplified debugging of Java stored procedures using the IBM VisualAge distributed debugger

There is a tutorial available through the Information Center that you can take to better understand this powerful new tool.

Memory Visualizer

This is a new tool in DB2 Version 8 that helps DBAs understand their systems' memory allocations on a per-instance basis. You can start this tool by right-clicking the instance you want to work with in the Control Center, and selecting **View Memory Usage**. You can also select the tool directly from the Start menu in a Windows environment. The following figure shows the Memory Visualizer monitoring an instance called DB2:



The top pane of the Memory Visualizer shows memory components organized in a navigation tree. Historical values, alarm, and warning thresholds are shown to the right of each component. The lower pane shows a memory usage plot.

The high-level memory components that this tool monitors includes database manager shared memory, database global memory, application global memory, agent/application shared memory, and agent private memory. Each high-level component is divided into lower-level components that determine how the memory is allocated and deallocated. For example, memory is allocated and deallocated when the database manager starts, when a database is activated, and when an application connects to a database.

Some of the key tasks that you can perform with the Memory Visualizer include:

- View overall memory usage
- Specify which memory information to display and which information to hide for a DB2 instance and its databases
- Update the configuration parameters for an individual memory component to prevent it from using too much or too little memory
- Save the memory allocation data

- Load memory allocation data from a file into a Memory Visualizer window
-

Storage Management

This is a new tool in DB2 Version 8 that helps DBAs understand their storage requirements and potential considerations. You can use this tool to schedule or immediately capture a snapshot of currently allocated storage resources. You can start the Storage Management tool by selecting a database in the Control Center, right-clicking, and selecting **Manage Storage**.

The Storage Management tool also allows you to specify warning and alarm thresholds (there are defaults provided out of the box) for the following storage criteria:

- **Space usage:** Measures the amount of disk space used in a database, database partition group, or tablespace
 - **Data skew:** Measures the balance of data between database partitions or tables
 - **Cluster ratio:** Measures the quality of coverage an index has on one table
-

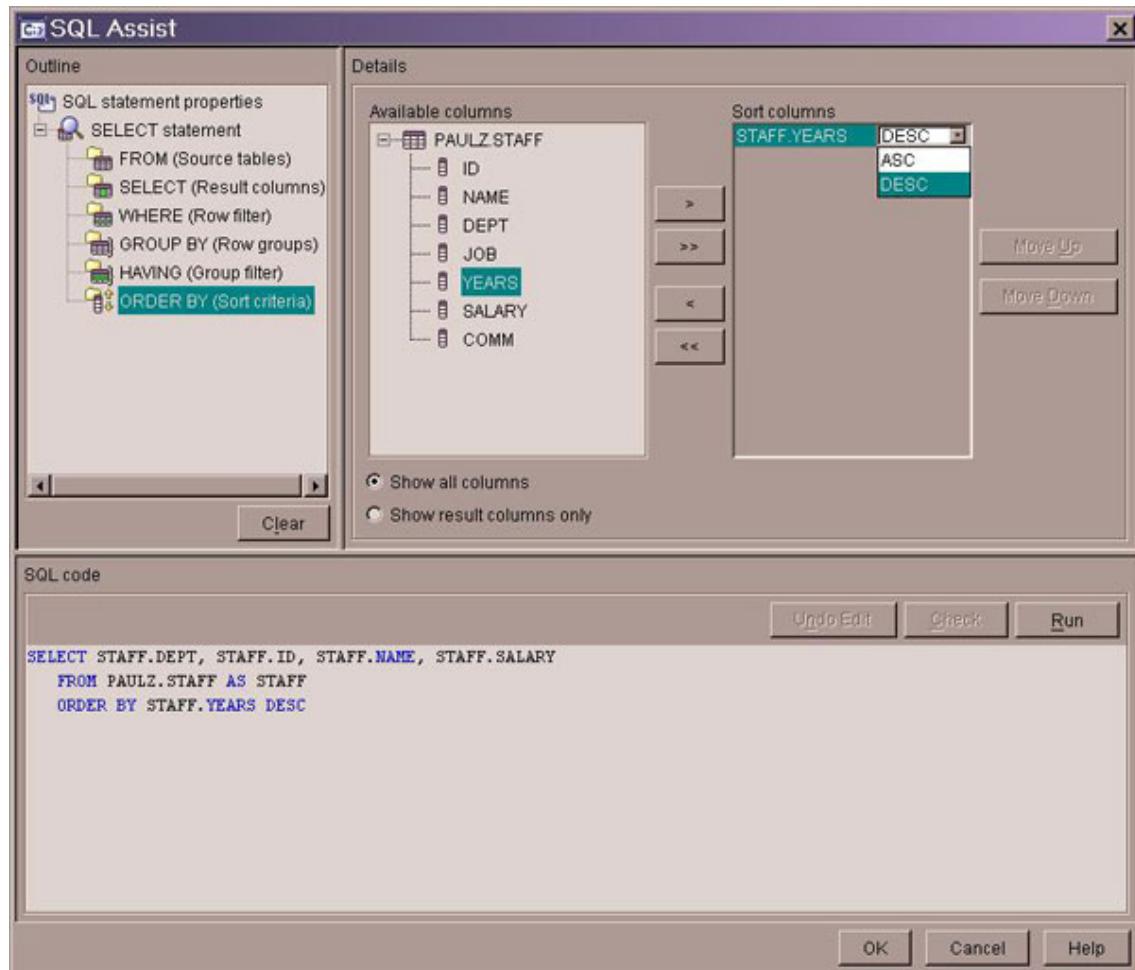
In-doubt Transaction Monitor

This is a new tool in DB2 Version 8. It helps DBAs work with global transactions that are in an in-doubt state. For example, a broken communication line could leave a transaction prepared but not yet committed or rolled back. In most cases, you won't have to use this tool and you can simply wait for the Transaction Manager to resynchronize; but if you cannot wait, you can use the In-doubt Transaction Manager. Only experienced DBAs should try this -- this tool is for people who *really* know what they are doing. You can start this tool by either selecting it from the Start menu, or entering the db2indbt command from a command prompt.

SQL Assist

SQL Assist is a graphical tool you can use to build SELECT, INSERT, UPDATE, and DELETE statements. The tool uses a notebook to help you organize the information you need to create an SQL statement. SQL Assist is available from multiple tools and wizards.

SQL Assist makes it really easy for novice users to point and click their way through query creation and build more complex SQL statements. The next time you query your database with the Command Center, or build a view, invoke the SQL Assist tool to get a feel for how much it can help you. The following figure shows a screen from the SQL Assist tool:



First Steps

First Steps is a graphical tool that helps get you started with DB2. First Steps has a number of options; all are available by clicking on the icon next to the desired action.

With First Steps, you can create sample databases, launch the DB2 tools for management and connectivity, work with the Business Intelligence tutorials, view the product library, find more DB2 resources on the Web, and take a quick tour of all the new features in DB2 Version 8. I suggest that you take the tutorial and the quick tour to prepare for the exam (trust me, they'll help you). In addition, if you haven't already done so, make sure you create the sample databases -- you can use them to reinforce what you learn in these tutorials.

Section 7. Data warehousing

What is data warehousing?

Systems that contain *operational data* -- the data that runs the daily transactions of a business -- contain information that business analysts can use to better understand how the business is operating. For example, they can see which products were sold in which regions at which time of year. This helps identify anomalies or to project future sales.

However, there are several problems if analysts access operational data directly:

- They might not have the expertise to query the operational database. For example, querying IMS databases requires an application program that uses a specialized type of data-manipulation language. In general, the programmers who have the expertise to query the operational database have a full-time job in maintaining the database and its applications.
- Performance is critical for many operational databases, such as databases for a bank. The system cannot handle users making ad-hoc queries on operational data stores. Imagine that you are doing your banking on the Internet and paying bills. When you hit the OK button, it usually takes only a few seconds to process a payment. Now, consider a bank analyst trying to figure out how to make more money from an existing customer base. The analyst runs a query which is so complex that your transaction now takes about 30 seconds to complete. Obviously that performance time is not acceptable (and neither are the new charges that the analyst is dreaming up). For this reason, operational data stores and reporting data stores (including OLAP databases) are generally separated.

However, over the last few years, reporting data stores have tended to become pseudo-operational and current. Such stores are called *operation data stores* (ODSs). Consider the telecommunications industry, for example. ODSs are popular with these companies, as they try to identify fraudulent charges as soon as possible. DB2 is one of the few databases that is well suited for both operational and reporting workloads.

- Operational data is not generally in the best format for use by business analysts. For example, sales data that is summarized by product, region, and season is much more useful to analysts than raw data.

Data warehousing solves these problems. In data warehousing, you create stores of *informational data* -- data that is extracted from operational data and then transformed and cleansed for end-user decision making. For example, a data warehousing tool might copy all the sales data from the operational database, perform calculations to summarize the data, and write the summarized data to a database that is separate from the operational data. End users can query the separate database (the *warehouse*) without affecting the operational databases.

DB2 products for data warehousing

DB2 is built for business intelligence. Not only is DB2 the world's most scalable database, but it also has a robust set of business intelligence functions. The Data Warehouse Center, included at no additional cost with DB2, provides SQL-based extract, transform, and load (ETL) capabilities to move and transform data. The Data Warehouse Center includes agents for Windows and AIX that transfer data directly between sources and targets.

The Data Warehouse Center is designed to handle basic data warehousing requirements in AIX and Windows environments. For full-powered data warehousing projects, IBM offers the DB2 Data Warehouse Manager (DB2 DWM).

IBM's DB2 DWM provides a distributed, heterogeneous infrastructure for designing, building, maintaining, governing, and accessing highly scalable and robust DB2 data warehouses. DB2 DWM significantly adds to the ETL capabilities of Data Warehouse Center, including an expanded set of supported operating systems (including Solaris, Linux, and AS/400), an end-user metadata management solution, a query tool, and a proactive query governing tool.

DB2 DWM supports a wide variety of relational and nonrelational data sources. You can fill your DB2 data warehouse with data from the most common relational databases (like any member of the DB2 family, Oracle, Sybase, Informix, and Microsoft SQL Server), flat files, nonrelational data stores, and over 40 other data sources. DB2 DWM adds over 100 built-in transformers that you can use to cleanse, restructure, correlate, standardize, and summarize data to make it more accurate and usable.

DB2 DWM adds the following features to DB2 for data warehousing:

- Increased scalability of data warehouses
- Fast deployment of data marts
- Comprehensive management and resource controls for DBAs
- Easy access to data and metadata by end users
- Enterprise reporting to develop and deliver reports to unlimited clients

The DB2 DWM is really a suite of four products:

- **Data Warehouse Center**, which comes with enhanced and expanded SQL transformations for ETL, more agents to support all DB2-supported operating systems, and more
- **Information Catalog Center**, an end-user level metadata repository
- **Query Patroller**, a powerful proactive query governor
- **Query Management Facility** (QMF) for Windows, a query reporting tool

We'll consider each of these in subsequent panels.

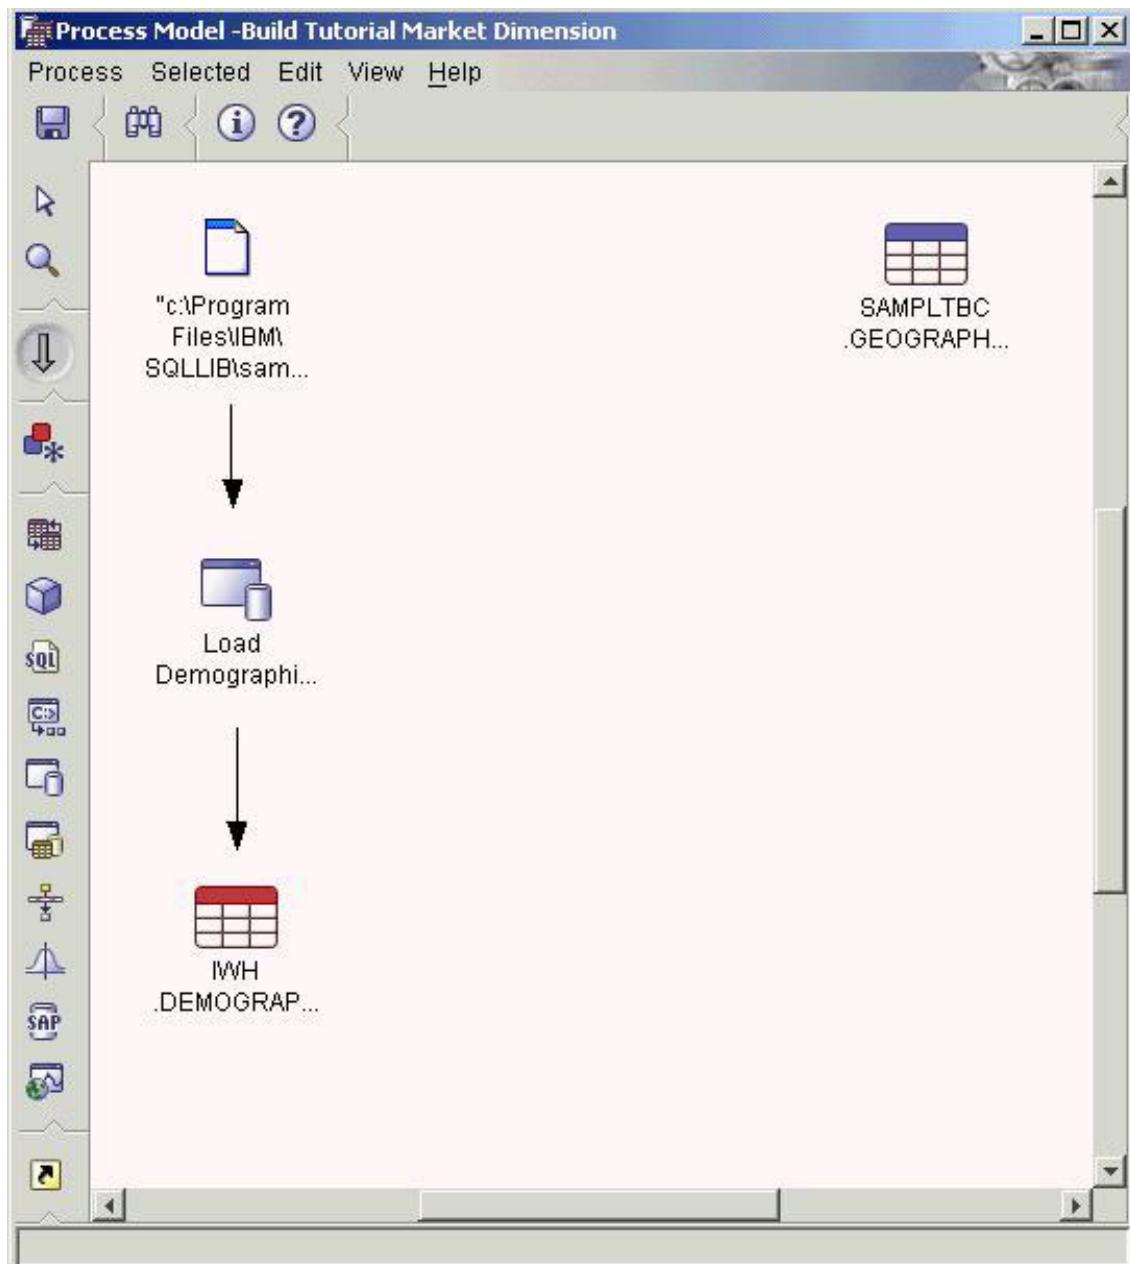
The Data Warehouse Center

DB2 on Windows and AIX comes with the Data Warehouse Center. This free tool offers a subset of the features of the Data Warehouse Center that is part of the DB2 DWM. The Data Warehouse Center that comes with DB2 is free of charge and allows DBAs a chance to sample some of the powerful features they can leverage with DB2 DWM.

The Data Warehouse Center offers:

- Registration and access to disparate data sources
- Definition of data extraction and transformation steps
- Population and the automation of warehouse management processes

The Data Warehouse Center supports full refresh and incremental update data movement options (including leveraging the power of IBM integrated data replication functions). The Warehouse Launchpad simplifies the job of populating a warehouse by leading you through related tasks. The figure below depicts the Data Warehouse Center's process modeling window:



The Data Warehouse Center is a component of a DB2 installation. If you select this component, the Control database will be created to help manage your data warehouse. A tutorial is included with DB2 that will take you through the steps of setting up and creating a data warehouse using this tool. I suggest you take it (you can find it in First Steps and in the DB2 Information Center).

Remember, the Data Warehouse Center is also part of the DB2 DWM. In the DB2 DWM, it comes with more agents, transforms, features, and so on, and is meant for enterprisewide data warehouses.

Information Catalog Center

The DB2 Information Catalog Center (DB2 ICC) is a tool that helps you find and access business information quickly and easily; it does for your organization what an electronic card catalog does for a library, and more.

In today's business environment, having access to important company information is a necessity. Gathering pertinent information is often a time-consuming, exhaustive process that involves searching for data in various locations and calling several departments or individuals for information.

The DB2 ICC streamlines the search for business information by providing a single information catalog and a robust searching capability for *metadata*. Metadata is data that describes data. For example, metadata might describe the source of data, or how it was assembled.

A company's information is maintained in different places: charts, graphs, displays, are in data warehouses; spreadsheets and tables of customer names, numbers, and other data are in accounting systems; text documents and graphic files are in office systems; and so on. All this data offers many opportunities for great timely decisions, but also presents a high degree of complexity.

You can catalog any or all of this information: The scope of what an information catalog contains depends on the information that can and should be shared within a company. The information cataloged doesn't actually have to reside in DB2 for you to take advantage of this tool. The DB2 ICC can be used to catalog information on a large mainframe system, a network server, your workstation, or even printed documents in a filing cabinet.

The DB2 ICC stores descriptive data about source information. This data can include the type of information, a description of the information, what it contains, who owns and updates it, and where and how to get to it. After you find the information you need, you can start spreadsheet programs, word processors, graphic tools, or other informational applications directly from the DB2 ICC. The DB2 ICC remains in the background while you work with the retrieved information in the application where the source information originates.

Effective metadata management is very important to efficient data access and analysis. The effective management of metadata adds value in several ways. It:

- Makes the integration of solutions easier, and frees skilled resources for other projects
- Improves impact analysis to rapidly respond to changing operational systems and changing business requirements
- Helps users find and understand the data they need for decision making

Query Patroller

DB2 Query Patroller (DB2 QP) provides DBAs with a facility for query governing and workload management. DB2 QP helps in the use and maintenance of any data

warehouse that is powered by DB2 on a UNIX or Windows operating system. Essentially, DB2 QP acts as a police force that can limit the resource consumption of a query, based on a user's profile. Management tools like this are becoming a necessity in data warehousing with the proliferation of ad-hoc query tools that allow you to ask business questions of that data (which isn't always the most efficient SQL).

DB2 QP was designed to address the needs of both users and administrators of reporting databases. Typically, each group has its own priorities and views concerning the database. For example, let's imagine an individual user named Tom. Tom may not pay attention to the workload caused by a query; because he isn't a SQL expert, he will often submit queries to a data warehouse that could be written more efficiently. But Tom is more interested in whether his report is ready or how he can schedule his query to run at a later date. Sometimes, because his result set has not been returned in a timely fashion, Tom might get the impression that his machine has crashed; he'll then restart the machine and resubmit the query. Now the data warehouse is going to be hit twice by the same query. Users like Tom need a way to gain insight into the progress of their queries to prevent such situations.

At the same time, a data warehouse DBA would be more interested in how they can give different users different priorities for query processing or in determining who is accessing what tables and at what time. The administrator needs a way to ensure that no user can take too much of the data warehouse's resources. The DBA also needs a way to manage queries, and to find out which tables are being accessed most frequently, in order to manipulate the physical architecture of the data warehouse so that these tables are stored on faster disks, for example.

DB2 Query Patroller's architecture has been completely redesigned in DB2 Version 8 and will not work with Version 7 servers. Despite the rearchitecture, the concept of this product is still the same.

You can combine DB2 Query Patroller features with other DB2 features, like the DB2 Governor, to provide a total query monitor environment.

Query Management Facility

Query Management Facility (QMF) for Windows is included with DB2 WM to provide a multipurpose query tool for business reporting, data sharing, server resource protection, robust application development, and native connectivity to all of the DB2 workstation platforms.

QMF is a tightly integrated, powerful, and reliable query and reporting tool set for DB2 databases on distributed and host platforms. It provides an environment that is easy for a novice to use, but also powerful enough for an application programmer. QMF offers extensive management and control over an enterprise query environment to restrict user access and protect valuable system resources.

In short, QMF also allows you to:

- Easily build queries and reports via a quick-start interface

- Leverage a Java-based query capability to launch queries from your favorite Web browser
- Integrate query results with desktop tools such as spreadsheets and personal databases
- Rapidly build data access and update applications
- Fully exploit DB2 performance, SQL syntax, and advanced database performance techniques such as static SQL

Section 8. OLAP

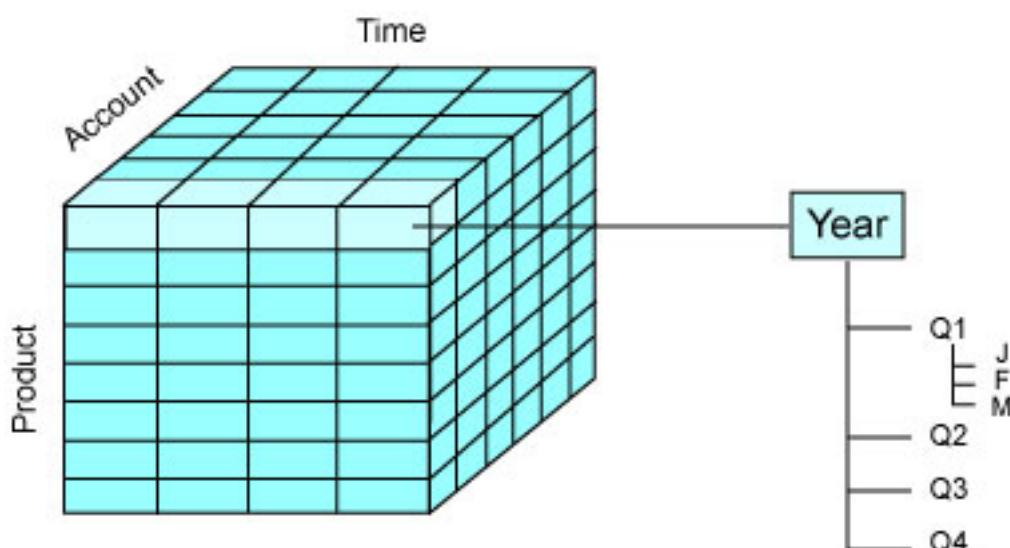
What is OLAP?

OLAP stands for Online Analytical Processing. Since the early 1990s, this multidimensional view of data has enhanced users' ability to leverage information in decision making. The technology is characterized by easy navigation between aggregated and derived information, delivering results with consistent, speed-of-thought response times.

Tools that employ OLAP technology, such as DB2, allow you to ask intuitive and complex ad-hoc questions about your business, such as "What is my profitability for the third quarter across the southeast region for my focus products?" Such a question requires multiple perspectives on data, such as time, regions, and products. These perspectives are called *dimensions*.

In OLAP analysis, users access data with numerous descriptors. Most business analysts are accustomed to viewing their business in two dimensions at a time. Think about it for a minute: How many of us have looked at a Lotus 1-2-3 spreadsheet? A spreadsheet represents facts in two dimensions. However, businesses usually don't operate in just two dimensions, but rather in multiple dimensions. As a result, we need a technology that can provide multidimensional analysis. (It should be noted that spreadsheets are often used as front ends to OLAP cubes, but are enhanced with the ability to leverage the multiple dimensions and drill down or through to other aggregations.)

OLAP analysis cannot be provided by a spreadsheet alone. For this reason, when we talk about multidimensional analysis, we talk about prepared *cubes* of information, as in the following figure:



The intersections of the axes in a cube are called *facts*, and they each represent a fact about the business. You can have hierarchies of facts in dimensions, and OLAP provides the ability to navigate these hierarchies. So if the Time dimension were

aggregated to quarters in the cube (as it is in the example cube above; as you can see, Time has four columns), you can imagine being able to look more closely at quarter one and see the members for the quarter: the months January, February, and March. This type of cube navigation is referred to as *drilling down*. Other navigation OLAP techniques include slicing, dicing, pivoting, or drilling up. Simply stated, OLAP can help your business make better, faster, and more informed decisions.

DB2 as an OLAP Server

There are several categories of OLAP, with the most common being multidimensional OLAP (*MOLAP*) and relational OLAP (*ROLAP*). DB2 can work with both. MOLAP is handled by the DB2 OLAP Server product and ROLAP by the native data stores in DB2.

An OLAP server processes multidimensional requests that calculate, consolidate, and retrieve information from a multidimensional database, a relational database, or both.

Some vendors store facts in proprietary formats, often called *multidimensional databases*, while others store them in a star-schema model in a relational database or cube. A sophisticated OLAP server offering like DB2 OLAP Server can access and store both multidimensional and relational databases.

MOLAP services are provided to customers through the DB2 OLAP server product. DB2 OLAP Server leverages the industry-leading Hyperion Essbase OLAP engine underneath its covers.

OLAP-like access to DB2 tables can be provided by SQL -- this is ROLAP. A variety of tools can generate this SQL, including DB2. DB2 includes the latest OLAP SQL (now part of SQL-99) including rank, rownumber, moving aggregates, covariance, and various linear regression functions.

The future of MOLAP is up for debate. Many users favor MOLAP because it is generally thought to be a faster OLAP approach for answering analytical questions. The drawback to MOLAP is that it requires a separate data store engine (the multidimensional engine), a separate administration approach, and a separate query language and API. DB2 OLAP Server (which is a repackaging of Hyperion's Essbase server) has its own storage engine that is separate from DB2, and also has its own APIs. Similarly, Microsoft's analysis services offering (their version of MOLAP, called MDX) has its own query language and management issues. Proponents of ROLAP feel that learning how to manage a new database engine and write a new query language is a burden, and prefer to leverage existing SQL skills to do OLAP. (As mentioned earlier, ROLAP is part of the SQL standard.) So the debate over the speed or cost of owning a totally different database and using a different query language rages on. It looks as if the industry will eventually head toward ROLAP. New approaches to generate better speed in ROLAP and an industry trend to cut costs are pushing customers into server and skill consolidation. Everyone has an opinion, and only time will tell. The truth is that both approaches will be around and be supported for years to come.

Section 9. Resources and feedback

Resources

Check out the other parts of the DB2 V8.1 Family Fundamentals Certification Prep tutorials:

- [*DB2 V8.1 Family Fundamentals Certification Prep, Part 2 of 6: DB2 Security*](#)
- [*DB2 V8.1 Family Fundamentals Certification Prep, Part 3 of 6: Accessing DB2 UDB Data*](#)
- [*DB2 V8.1 Family Fundamentals Certification Prep, Part 4 of 6: Working with DB2 UDB Data*](#)
- [*DB2 V8.1 Family Fundamentals Certification Prep, Part 5 of 6: Working with DB2 UDB Objects*](#)
- [*DB2 V8.1 Family Fundamentals Certification Prep, Part 6 of 6: Data Concurrency*](#)

For more information on the available DB2 products:

- Read "[*Which Distributed Edition of DB2 Version 8 is Right for You?*](#)" (by Paul Zikopoulos and Roman Melnyk) on DB2 Developer Domain.
- Review the [*Distributed DB2 Version 8*](#) Web site.
- Download [*DB2 Personal Developer's Edition*](#)

Learn more about some of the products mentioned in this tutorial:

- [*DB2 Extenders*](#)
- [*DB2 Data Links Manager*](#)
- [*DB2 Connect*](#)
- [*IBM Data Management Tools*](#)

For more information on basic tools functions and wizards:

- Play with the Control Center and expand each folder and select the options from the pop-up menus. Online help is available with each wizard that will not only explain the reason for the using the wizard, but what each of the fields mean as well.
- Check out the [*Guide to GUI Tools for Administration and Development*](#).
- Start and experiment with the Configuration Assistant. Use it to add a new database to your system and configure it. Refer to this tool's online help for more information.
- Take the new Development Center tutorial, available from the Information Center.
- Read "[*DB2 Development Center: The Next Generation AD Tooling for DB2*](#) by Abdul Al-Azzawe, also on the DB2 Developer Domain.

For more information on the DB2 processors:

- Refer to Chapter 2, "Command Line Processor (CLP)," in the [*DB2 Command Reference*](#).

- Enter the examples I walked through in the readings.

You can learn more about data warehousing and OLAP from the following resources:

- The Business Intelligence Tutorial, which is available from First Steps or the Information Center.
- The Extended Business Intelligence Tutorial, which is available from the Information Center (it requires a download that will be launched from the Information Center).
- The DB2 Quick Tour that you can launch from the launchpad when performing a DB2 installation or from First Steps.
- The [DB2 Warehouse Manager](#) Web site.
- The [DB2 OLAP Server](#) Web site.
- The [Business Intelligence solutions](#) Web site.
- Read "[DB2 Query Patroller 101](#)" by Paul Zikopoulos on IDUG's site.

For more information on the DB2 Fundamentals Exam 700:

- [IBM Data Management Skills information](#)
 - Download a [self-study course for experienced Database Administrators \(DBAs\)](#) to quickly and easily gain skills in DB2 UDB.
 - Download a [self study course for experienced relational database programmers](#) who'd like to know more about DB2.
 - [General Certification information](#) - including some book suggestions, exam objectives, courses.
-

Feedback

Colophon

This tutorial was written entirely in XML, using the developerWorks Toot-O-Matic tutorial generator. The open source Toot-O-Matic tool is an XSLT style sheet and several XSLT extension functions that convert an XML file into a number of HTML pages, a zip file, JPEG heading graphics, and two PDF files. Our ability to generate multiple text and binary formats from a single source file illustrates the power and flexibility of XML. (It also saves our production team a great deal of time and effort.)

You can get the source code for the Toot-O-Matic at www6.software.ibm.com/dl/devworks/dw-tootomatic-p. The tutorial [Building tutorials with the Toot-O-Matic](#) demonstrates how to use the Toot-O-Matic to create your own tutorials. developerWorks also hosts a forum devoted to the Toot-O-Matic; it's available at www-105.ibm.com/developerworks/xml_df.nsf/AllViewTemplate?OpenForm&RestrictToCategory=11. We'd love to know what you think about the tool.